


Manage data

Qlik Sense®
April 2019
Copyright © 1993-2019 QlikTech International AB. All rights reserved.



© 2019 QlikTech International AB. All rights reserved. Qlik®, Qlik Sense®, QlikView®, QlikTech®, Qlik Cloud®, Qlik DataMarket®, Qlik Analytics Platform®, Qlik NPrinting®, Qlik Connectors®, Qlik GeoAnalytics®, Qlik Core®, Associative Difference®, Lead with Data™, Qlik Data Catalyst™, Qlik Associative Big Data Index™ and the QlikTech logos are trademarks of QlikTech International AB that have been registered in one or more countries. Other marks and logos mentioned herein are trademarks or registered trademarks of their respective owners.

1 About this document	9
2 Managing data	10
3 Managing data in the app with Data manager	11
3.1 Previewing a data table	11
3.2 Adding a new data table	11
3.3 Editing a data table	12
3.4 Deleting a data table	12
3.5 Managing data table associations	12
3.6 Applying changes and reloading data	13
3.7 Undo and Redo actions in Data manager	14
3.8 Viewing table transformation details in Data manager	14
3.9 Interaction between Data manager and the data load script	14
3.10 Concatenating tables in Data manager	15
3.11 Adding data to the app	15
In-App	16
File locations	16
Data connections	16
Data content	17
Attach files to this app	17
Connect to a new data source	17
Add data	17
Which data sources are available to me?	17
Adding data from an existing data source	17
Adding data from a new data source	19
Adding data manually in Qlik Sense	21
Attaching data files and adding the data to the app	23
Selecting data fields	26
Filtering data from files	31
Filtering data from data connectors	33
Adding data from Qlik DataMarket	33
3.12 Editing a table	46
Renaming a table	46
Renaming a field	46
Managing associations to other tables	47
Changing field type and display format	47
Hiding fields from analysis	47
Assessing table field data before loading data	48
Replacing field values in a table	48
Setting field values as null in a table	48
Setting a custom order for field values	48
Splitting a field in a table	48
Grouping measure data into ranges	49
Viewing field transformation details	49
Unpivoting crosstab data	49

Updating a table from the data source	49
Adding a calculated field	50
Sorting a table	50
Undo and redo actions	50
Associating data in the table editor	50
Using calculated fields	52
Changing field types	65
Hiding fields from analysis	67
Assessing table field data before loading data	68
Replacing field values in a table	69
Setting field values as null in a table	71
Customizing the order of dimension values	72
Splitting a field in a table	73
Grouping measure data into ranges	75
Unpivoting crosstab data in the data manager	78
3.13 Concatenating tables in Data manager	80
Automatically concatenating tables	80
Forcing concatenation between tables	81
Splitting concatenated tables	83
3.14 Viewing table and field transformation details in Data manager	84
Viewing table details	84
Viewing field details	84
3.15 Step-by-step - Combining tables using forced concatenation	85
Concatenation at a glance	85
Walkthrough - Forced concatenation	85
A step further - adding a new table and concatenating the data fields	92
3.16 Synchronizing scripted tables in Data manager	94
Synchronizing scripted tables	95
Removing managed scripted tables	95
3.17 Managing data associations	95
Associating tables using the Recommended associations panel	96
Associating tables manually	98
Breaking associations	99
Editing associations	99
Previewing data	100
Synthetic keys	100
Limitations	100
Applying changes and reloading data	101
4 Loading data with the data load script	102
4.1 Interaction between Data manager and the data load script	102
4.2 Using the data load editor	103
A: Toolbar	103
B: Data connections	103
C: Text editor	104
D: Sections	104

E: Output	104
Quick start	104
Toolbars	104
Connect to data sources in the data load editor	106
Select data in the data load editor	108
Edit the data load script	114
Organizing the script code	117
Debug the data load script	118
Saving the load script	121
Run the script to load data	121
Keyboard shortcuts in the Data load editor	122
4.3 Understanding script syntax and data structures	123
Extract, transform and load	123
Data loading statements	124
Execution of the script	125
Fields	125
Logical tables	130
Data types in Qlik Sense	140
Dollar-sign expansions	143
Using quotation marks in the script	146
Wild cards in the data	149
NULL value handling	151
4.4 Guidelines for data and fields	153
Guidelines for amount of loaded data	154
Upper limits for data tables and fields	154
Recommended limit for load script sections	154
Conventions for number and time formats	154
4.5 Working with QVD files	158
Purpose of QVD files	158
Creating QVD files	159
Reading data from QVD files	159
QVD format	159
4.6 Managing security with section access	160
Sections in the script	160
Dynamic data reduction	162
Inherited access restrictions	164
4.7 Configuring analytic connections in Qlik Sense Desktop	164
Qlik open source SSE repositories	165
Description of the elements	165
5 Managing big data with on-demand apps	167
5.1 On-demand app components	167
5.2 Constructing on-demand apps	168
5.3 Publishing on-demand apps (Windows)	169
5.4 Sharing on-demand apps (Kubernetes)	169
5.5 Advantages of on-demand apps	169

5.6	Creating an on-demand selection app	170
5.7	Creating an on-demand template app	171
	Structure of a template app	171
	Single Sign-On (SSO)	172
	Reload nodes for template apps	173
	Binding expressions in on-demand template apps	173
5.8	Building an on-demand app	177
6	Connecting to data sources	181
6.1	Create a connection	181
6.2	Data connection types	181
	Attached files	181
	Qlik DataMarket	181
	Database connectors	182
	Essbase	182
	Local or network files	182
	ODBC connections through DSN	182
	Qlik Associative Big Data Index	183
	REST	183
	Salesforce	183
	SAP	183
	Web files	183
	Web Storage Provider Connectors	183
	Third-party connectors	183
6.3	Where is the data connection stored?	184
6.4	Limitations	184
6.5	Loading data from files	184
	How can I load data from files?	184
	Loading files from local and network file folders	185
	Loading files from web resources	185
	Loading data from Microsoft Excel spreadsheets	187
6.6	Loading data from databases	189
	Loading data from an ODBC database	189
	ODBC	190
	OLE DB	193
	Logic in databases	195
6.7	Loading data from Qlik DataMarket	195
6.8	Accessing large data sets with Direct Discovery	196
	Differences between Direct Discovery and in-memory data	197
	Direct Discovery field types	202
	Data sources supported in Direct Discovery	203
	Limitations when using Direct Discovery	205
	Multi-table support in Direct Discovery	207
	Using subqueries with Direct Discovery	208
	Logging Direct Discovery access	211

7 Viewing the data model	212
7.1 Toolbar	212
7.2 Moving and resizing tables in the data model viewer	213
Moving tables	213
Resizing tables	214
7.3 Preview of tables and fields in the data model viewer	214
Showing a preview of a table	214
Showing a preview of a field	215
7.4 Creating a master dimension from the data model viewer	215
7.5 Creating a master measure from the data model viewer	216
8 Best practices for data modeling	217
8.1 Turning data columns into rows	217
8.2 Turning data rows into fields	217
8.3 Loading data that is organized in hierarchical levels, for example an organization scheme	218
8.4 Loading only new or updated records from a large database	219
8.5 Combining data from two tables with a common field	219
8.6 Matching a discrete value to an interval	219
8.7 Handling inconsistent field values	220
8.8 Handling inconsistent field value capitalization	221
8.9 Loading geospatial data to visualize data with a map	222
8.10 Loading new and updated records with incremental load	222
Append only	223
Insert only (no update or delete)	223
Insert and update (no delete)	224
Insert, update and delete	225
8.11 Combining tables with Join and Keep	226
Joins within a SQL SELECT statement	226
Join	227
Keep	227
Inner	227
Left	228
Right	230
8.12 Using mapping as an alternative to joining	231
8.13 Working with crosstables in the data load script	233
Unpivoting a crosstab with one qualifying column	233
Unpivoting a crosstab with two qualifying columns	234
8.14 Generic databases	235
8.15 Matching intervals to discrete data	237
Intervalmatch example	237
Using the extended intervalmatch syntax to resolve slowly changing dimension problems	238
8.16 Creating a date interval from a single date	240
8.17 Loading hierarchy data	243
8.18 Loading your own map data	244

Supported name data for fields in a map visualization	244
Loading point and area data from a KML file	245
Loading map data with data profiling	245
Loading and formatting point data	246
8.19 Data cleansing	248
Mapping tables	248
Using a mapping table	249
9 Troubleshooting - Loading data	251
9.1 A data connection stops working after SQL Server is restarted	251
9.2 Synthetic keys warning when loading data	251
9.3 Circular references warning when loading data	251
9.4 Data selection problems with an OLE DB data source	252
9.5 Character set problems with non-ANSI encoded data files	252
9.6 Connector is not working	252
The connector is not properly installed	253
The connector is not adapted for Qlik Sense	253
9.7 Data load script is executed without error, but data is not loaded	253
A statement is not terminated with a semicolon	253
Single quote character inside a string	253
9.8 Columns are not lining up as expected when selecting data from a fixed record file	254
9.9 Error message "Invalid path" when attaching a file	254
9.10 Errors when loading an app converted from a QlikView document	254
Absolute file path references are used in the script	255
Unsupported functions or statements are used in the script	255
9.11 Problems connecting to, and loading data from Microsoft Excel files through ODBC	255
9.12 Attaching a file by dropping it in Add data does not work	255
9.13 Tables with common fields are not automatically associated by field name	256
9.14 Date fields are not recognized as date fields in sheet view	256
Data profiling was disabled when the table was added	256
Date format was not recognized	257
9.15 Data manager does not show tables in app that contains data	257
9.16 Data manager workflows are broken for all users creating apps on a server	257
9.17 Running out of disk space	258
9.18 Data load editor does not display the script	258

1 About this document

This document describes how to add and manage data, how to build a data load script for more advanced data models, how to view the resulting data model in the data model viewer, and presents best practices for data modeling in Qlik Sense.



For detailed reference regarding script functions and chart functions, see the [Script syntax and chart functions](#).

This document is derived from the online help for Qlik Sense. It is intended for those who want to read parts of the help offline or print pages easily, and does not include any additional information compared with the online help.

You find the online help, additional guides and much more at help.qlik.com/sense.

2 Managing data

When you have created a Qlik Sense app, the first step is to add some data that you can explore and analyze. This section describes how to add and manage data, how to build a data load script for more advanced data models, how to view the resulting data model in the data model viewer, and presents best practices for data modeling in Qlik Sense.

There are two ways to add data to the app.

- **Data manager**

You can add data from your own data sources, or from other sources such as Qlik DataMarket, without learning a script language. Data selections can be edited, and you can get assistance with creating data associations in your data model.

- **Data load editor**

You can build a data model with ETL (Extract, Transform & Load) processes using the Qlik Sense data load script language. The script language is powerful and enables you to perform complex transformations and creating a scalable data model.





*You can convert a data model built in **Data manager** into a data load script, which can be developed further in **Data load editor**, but it is not possible to convert a data load script to a **Data manager** data model. The **Data manager** data model and data tables defined in the data load script can still co-exist, but this can make it harder to troubleshoot problems with the data model.*

3 Managing data in the app with Data manager

Add and manage data from the **Data manager** so that you can use the data in your app.

There are two views in data manager:

-  **Associations**
You can create and edit association between tables.
-  **Tables**
You get an overview of all data tables in the app, whether you added them using **Add data**, or loaded them with the data load script. Each table is displayed with the table name, the number of data fields, and the name of the data source.

3.1 Previewing a data table


You can preview a table to see which columns it contains, and a sample set of the data.

Do the following:

- Select the data table you want to preview.

A preview of the table data set is displayed.

3.2 Adding a new data table

You can quickly add a data table to your app. Open the **Data manager** and then click **+**. You can also click **Add data** in the  menu. You are also prompted to add data when you create a new app.

You can add data from the following data sources:


Data sources

Data source	Description
In-App	Select from data sources that are available in your app. These can be files that you have attached to your app. You can also create a data source and manually add data to it using Manual entry .
File locations	Select from files on a network drive, for example a drive that has defined by your administrator.
Data connections	Select from existing data connections that have been defined by you or an administrator.
Data content	Select from Qlik DataMarket normalized data from public and commercial databases.

3.3 Editing a data table


You can edit all the data tables that you have added with **Add data**. You can rename the table and fields in the data table, and update the fields from the data source. It is also possible to add a calculated field and adjust date and time formats.

Do the following:

1. Click  on the data table you want to edit.
The data table editor opens, and you can perform the edits and transformations you want to do.
2. Click **Close** to return.

The table is now marked **Pending update**, and the changes will be applied to the app data the next time you reload data.




*You can only edit data tables added with **Add data**. If you click  on a table that was loaded using the load script, the data load editor opens. For more information, see [Using the data load editor](#) (page 103).*

3.4 Deleting a data table

You can only delete data tables added with **Add data**. Data tables that were loaded using the load script can only be removed by editing the script in the data load editor.

Do the following:

- Click  on the data table you want to delete.

The table is now marked **Pending delete** and will be removed the next time you reload data.

You can undo and redo your delete actions by clicking  and .



If you have used fields from the data table in a visualization, removing the data table will result in an error being shown in the app.

3.5 Managing data table associations

When you add several tables that need to be associated, the perfect situation is that the tables associate with key fields that have identical names in the different tables. If that is the case, you can add them to Qlik Sense with the data profiling disabled option of **Add data**, and the result will be a data structure with the tables associated correctly.

If you have less than ideal data sources, there are a number of possible association problems.


3 Managing data in the app with Data manager

- If you have loaded two fields containing the same data but with a different field name from two different tables, it's probably a good idea to name the fields identically to relate the tables.
- If you have loaded two fields containing different data but with identical field names from two different tables, you need to rename at least one of the fields to load them as separate fields.
- If you have loaded two tables containing more than one common field.

If you want to associate your data, we recommend that you use the **Add data** option with data profiling enabled. This is the default option. You can verify this setting by clicking **•••** beside the **Add data** button in the lower right corner of the Add Data page.

Qlik Sense performs data profiling of the data you want to load to assist you in fixing the table association. Existing bad associations and potential good associations are highlighted, and you get assistance with selecting fields to associate, based on analysis of the data.

You can manage table associations in two different ways:

- In the  **Associations** view of data manager.
You can create associations based on recommendations, or create custom associations based on one or more fields.
- Using the **Associate** option in the table editor.
You can create custom associations and composite key associations based on several fields.




If you disable data profiling when adding data, Qlik Sense will associate tables based on common field names automatically.

3.6 Applying changes and reloading data

Changes that you have made in the **Data manager** will not be available in the app until you have reloaded data. When you reload data, the changes are applied and any new data that you have added is loaded from the external data sources. Data that you loaded previously is not reloaded.




You can reload all the data from the external data sources by using the  button in the **Data manager** footer.



The  button reloads all the data for the selected table. It does not reload all the data for all the tables in the app.

If the data in **Data manager** is out of sync with the app data, the **Load data** button is green. In the **Associations** view, all new or updated tables are indicated with *, and deleted tables are a lighter color of gray. In the **Tables** view, all new, updated, or deleted tables are highlighted in blue and display an icon that shows the status of the table:

3 Managing data in the app with Data manager

- Tables marked with **Pending delete**  will be deleted.
- Tables marked with **Pending update**  will be updated with fields that have been added, renamed, or removed, or the table will be renamed.
- Tables marked with **Pending add**  will be added.


Do the following:

- Click **Load data** to load the changes in the app.



The app data is now updated with changes you made in **Data manager**.

To apply changes and reload all the data in the selected table from the external data sources:

Do the following:

- Click the  button in the **Data manager** footer.

3.7 Undo and Redo actions in Data manager

When you are editing in **Data manager**, you can undo or redo some actions by clicking  and , or by using the keyboard shortcuts Ctrl + Z and Ctrl + Y.

The log of actions is cleared if you:

- Change view, for example, going from the table overview to **Associations**.
- Load data.
- Close **Data manager**.

3.8 Viewing table transformation details in Data manager

You can view the operations and transformations performed on a table in **Data manager** using the **Details** dialog. The **Details** dialog is available in the **Associations** and **Table** views.

Details displays the current operations and transformations made to the selected table. This shows you the source of a table, the current changes that have been made, and the sequence in which the changes have been applied. **Details** enables you to more easily understand how a table got into a current state. You can use **Details**, for example, to easily see the order in which tables were concatenated.

3.9 Interaction between Data manager and the data load script

When you add data tables in the **Data manager**, data load script code is generated. You can view the script code in the **Auto-generated section** of the data load editor. You can also choose to unlock and edit the generated script code, but if you do, the data tables will no longer be managed in the **Data manager**.

3 Managing data in the app with Data manager

By default, data tables defined in the load script are not managed in **Data manager**. That is, you can see the tables in the data overview, but you cannot delete or edit the tables in **Data manager**, and association recommendations are not provided for tables loaded with the script. If you synchronize your scripted tables with **Data manager**, however, your scripted tables are added as managed scripted tables to **Data manager**.



*If you have synchronized tables, you should not make changes in the data load editor with **Data manager** open in another tab.*


You can add script sections and develop code that enhances and interacts with the data model created in **Data manager**, but there are some areas where you need to be careful. The script code you write can interfere with the **Data manager** data model, and create problems in some cases, for example:

- Renaming or dropping tables added with **Data manager** in the script.
- Dropping fields from tables added with **Data manager**.
- Concatenation between tables added with **Data manager** and tables loaded in the script.
- Using the **Qualify** statement with fields in tables added with **Data manager**.
- Loading tables added with **Data manager** using **Resident** in the script.
- Adding script code after the generated code section. The resulting changes in the data model are not reflected in **Data manager**.

3.10 Concatenating tables in Data manager

Concatenation combines two tables into a single table with combined fields. It consolidates content, reducing the number of separate tables and fields that share content. Tables in **Data manager** can be automatically or forcibly concatenated.

Concatenated table in Tables view and Associations view.



3.11 Adding data to the app

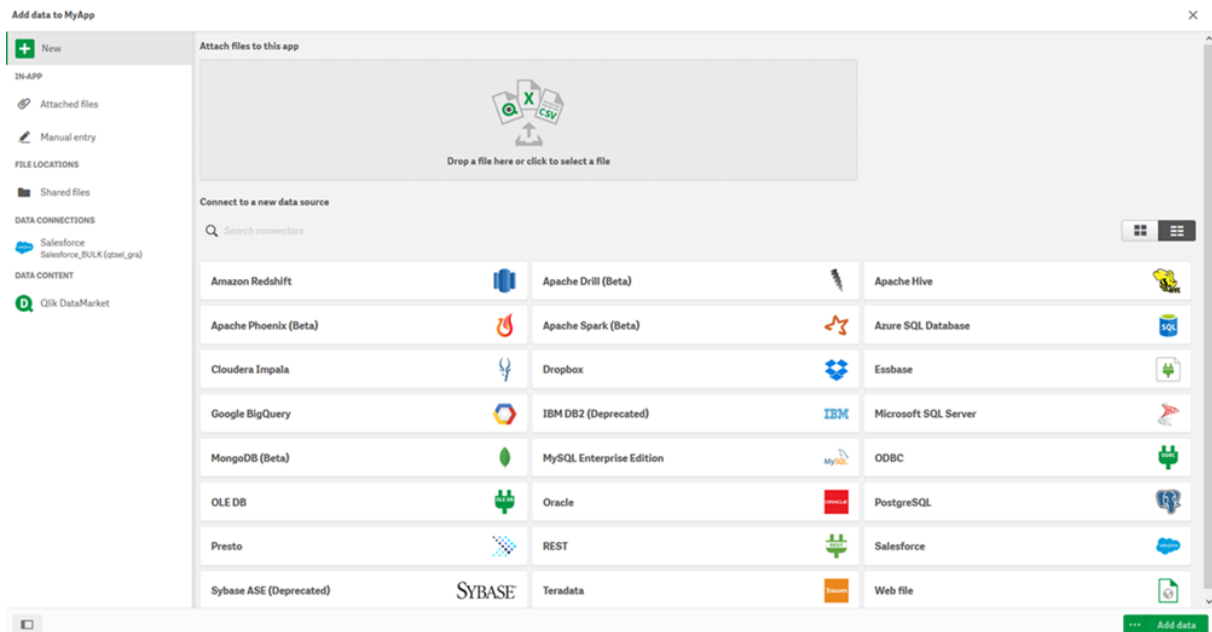
You can add data to your app quickly. Open the **Data manager** and then click **+**. You can also click **Add data** in the **☰** menu. You are also prompted to add data when you create a new app. When you are editing a sheet, you can also click **Add** in the **Fields** panel to add data.

3 Managing data in the app with Data manager



The add data options and data sources that are available to you depend on your Qlik Sense platform and configuration.

Add data view in Qlik Sense Enterprise on Windows



In-App

Attached files. Platforms: Qlik Sense Enterprise on Windows, Qlik Sense Cloud Business. Click to view the files that are attached to the app. You can load data from these files.

Manual entry. Platforms: All. Click to create a table in-app and add it to **Data manager**.

File locations

Data Files. Platforms: Qlik Cloud Services, Kubernetes, Qlik Sense Cloud Business.

Shared data files. Platforms: Qlik Sense Enterprise on Windows. This folder appears if your administrator has defined a network folder that contains shared files.

My computer. Platforms: Qlik Sense Desktop

Click to upload a data file, or to add data from a file that has already been uploaded.

Data connections

Platforms: All.

Displays connections that have been created to an external data source. The connection appears after it has been created under **Connect to a new data source**.

Click a connection to add data to the app.

Data content

Qlik DataMarket. Platforms: Qlik Sense Enterprise on Windows, Qlik Sense Cloud Business, Qlik Sense Desktop.

The Qlik DataMarket sources that are available to you depend on your subscription.

Click to add data to the app.

Attach files to this app

Click to attach a file to the app.

Connect to a new data source

Click to create a connection to a new data source.

Add data

Click to add data to an app. The button is enabled after you have created a connection and selected your data to load. You can add data with profiling enabled or disabled. Data profiling is recommended and enabled by default. Click ●●● to disable data profiling.

Which data sources are available to me?

What data source types are available to you depends on a number of factors:

- Access settings
Administrator settings determine which types of data sources you can connect to.
- Installed connectors
Qlik Sense contains built-in support for many data sources. Built-in connectors are installed automatically by Qlik Sense. To connect to additional data sources you may need to install connectors separately that are specifically for those data sources. Such separately installed connectors are supplied by Qlik or a third party.
- Local file availability
Local files on your desktop computer are only available in Qlik Sense Desktop. They are not available for use with a server installation of Qlik Sense.



If you have local files that you want to load on a server installation of Qlik Sense, you need to attach the files to the app, or transfer the files to a folder available to the Qlik Sense server, preferably a folder that is already defined as a folder data connection.


Adding data from an existing data source

You can add data to your app from connections that have already been defined by you or an administrator. These can be a database, a folder containing data files, or a connector to an external data source, such as Salesforce.com.

3 Managing data in the app with Data manager





Do not add a table in **Data manager** that has already been added as a scripted table with the same name and same columns in **Data load editor**.

You can delete connections from **Add data** by right-clicking the connection and selecting  **Delete connection**.



If you delete a connection, you must delete any tables from **Data manager** that used that connection before you load data.

Do the following:

1. Open your app.
2. Open the **Data manager** and then click . You can also click **Add data** in the  menu.
3. Under **Data connections**, select an existing connection.
Some connections go directly to their data sources, where you select tables and fields to load. For example, connections to Salesforce.com or a database, such as IBM DB2, go directly to the source for data selection.
4. Select the specific data source you want to add data from if the connection offers a selection. This differs depending on the type of data source.
 - File-based data sources: Select a file.
 - Databases: Set which database to use.
 - Web files: Enter the URL of the web file.
 - Other data sources: Specified by the connector.
5. Select the tables and fields to load.
6. Optionally, select to apply a data filter if you want to select a subset of the data contained in the fields you have selected.

If your data source is a file, select **Filters**. Beside the table to which you want to add a filter, click **Add filter**, select a field, select a condition, and then enter a value with which to filter.



Qlik Sense does not support filters on date fields from QVD files.

Note the following:

- You can apply multiple filters to the same field.
- You can remove filters in the **Associations** view of **Data manager** or from **Select data from source**. For the changes to take effect, reload data by clicking the **Load data** button.

For databases and connectors, when you select **Filter data**, a text box opens for the filter criteria.

Note the following:

- Filters are applied to field names from the database. If you rename a field in the **Data manager**, you have to apply the filter to the original field name from the database. For example, if a field is named EMP in your database, and you rename it to EMPLOYEE in the **Data manager**, you have to

3 Managing data in the app with Data manager

apply the filter EMP = 'filter_value'.

- You can clear data filters in the **Associations** view of the **Data manager**. For the changes to take effect, reload data by clicking the **Load data** button. You have to split concatenated tables before clearing filters.
 - Filtering data is not currently available for all connectors, or for attached files. Filtering is also not available for Qlik DataMarket.
7. Click **Add data** to open the data in the **Associations** view of the data manager. This allows you to continue to add data sources, transform the data, and associate the tables in **Data manager**. Data profiling is enabled by default when you click **Add data**. Data profiling does the following:
- Recommends data associations.
 - Auto-qualifies common fields between tables. This adds a unique prefix based on table name.
 - Maps date and time fields to autoCalendar.


Tables are not associated on common field names automatically. You can associate tables in the **Associations** view.



*If you want to load the data directly into your app, click **...** and then disable data profiling. This will load the newly selected data from the external data source when you add data. Tables will be associated on common field names automatically. Date and time fields will not be created.*

For more information, see *Managing data associations (page 95)*.

8. Click **Load data** when you are done preparing the data. If serious problems are detected, you need to resolve the problems in **Data manager** before you can load data into the app.
- For more information, see *Troubleshooting - Loading data (page 251)*.


To reload all the data that you have selected from the external source, use the  button in the **Data manager** footer. This ensures you get all the current data from the source for the selections you have made. Reloading all the data can take longer than loading only the new data. If the data you loaded previously has not been changed in the data source, it is not necessary to reload all the data.

Adding data from a new data source

You can add data to your app from a new data source. When you add data from a new data source, a connection to the data source is created in Data Connections, making it easier to add more data from the same data source.



*Do not add a table in **Data manager** that has already been added as a scripted table with the same name and same columns in **Data load editor**.*

You can delete connections from **Add data** by right-clicking the connection and selecting  **Delete connection**.

3 Managing data in the app with Data manager



If you delete a connection, you must delete any tables from **Data manager** that used that connection before you load data.

Do the following:

1. Open an app.
2. Open the **Data manager** and then click **+**. You can also click **Add data** in the **☰** menu.
3. Under **Connect to a new data source**, select a source.
4. Enter the connection parameters required by the data source.
For example:
 - File based data sources require that you specify a path to the files and select a file type.
 - Databases such as Oracle and IBM DB2 require database properties and access credentials.
 - Web files require the URL of the web file.
 - ODBC connections require DSN credentials.
5. Select the tables and fields to load.
6. Optionally, select to apply a data filter if you want to select a subset of the data contained in the fields you have selected.
If your data source is a file, select **Filters**. Beside the table to which you want to add a filter, click **Add filter**, select a field, select a condition, and then enter a value with which to filter.



Qlik Sense does not support filters on date fields from QVD files.

Note the following:

- You can apply multiple filters to the same field.
- You can remove filters in the **Associations** view of **Data manager** or from **Select data from source**. For the changes to take effect, reload data by clicking the **Load data** button.

For databases and connectors, when you select **Filter data**, a text box opens for the filter criteria.

Note the following:

- Filters are applied to field names from the database. If you rename a field in the **Data manager**, you have to apply the filter to the original field name from the database. For example, if a field is named EMP in your database, and you rename it to EMPLOYEE in the **Data manager**, you have to apply the filter EMP = 'filter_value'.
 - You can clear data filters in the **Associations** view of the **Data manager**. For the changes to take effect, reload data by clicking the **Load data** button. You have to split concatenated tables before clearing filters.
 - Filtering data is not currently available for all connectors, or for attached files. Filtering is also not available for Qlik DataMarket.
7. Click **Add data** to open the data in the **Associations** view of the data manager. This allows you to continue to add data sources, transform the data, and associate the tables in **Data manager**. Data profiling is enabled by default when you click **Add data**. Data profiling does the following:

3 Managing data in the app with Data manager

- Recommends data associations.
- Auto-qualifies common fields between tables. This adds a unique prefix based on table name.
- Maps date and time fields to autoCalendar.

Tables are not associated on common field names automatically. You can associate tables in the **Associations** view.



*If you want to load the data directly into your app, click **•••** and then disable data profiling. This will also reload all existing data from data sources when you add the data. Tables will be associated on common field names automatically. Date and time fields will not be created.*

For more information, see *Managing data associations (page 95)*.

8. Click **Load data** when you are done preparing the data. If serious problems are detected, you need to resolve the problems in **Data manager** before you can load data into the app.

For more information, see *Troubleshooting - Loading data (page 251)*.

Adding data manually in Qlik Sense

Manual entry in **Add data** enables you to enter data into an editor and add it as a table in **Data manager**.

Manually entering data is useful if you want to use a limited set of data from another source. For example, if you only wanted a selection of rows from an Excel spreadsheet or from a table on a webpage to be loaded as a table into **Data manager**, you could copy and paste the selected data into **Manual entry** and then add it as a table in **Data manager**. It is also useful if you have a small amount of data that might be faster to add manually than importing from another data source.

To add data manually, you open **Add data**, select **Manual entry**, enter your data into the table, and then add the table to **Data manager**. The table editor starts with one row and two columns, but as you add data to the table, additional empty rows and columns are automatically added to the table.



***Manual entry** does not automatically save as data is entered. Data entered may be lost if the screen is refreshed, if the session times out, or if the connection is lost before the data is added to **Data manager**.*

In addition to typing data, you can copy and paste it from other sources. **Manual entry** preserves the columns and rows of data copied from Excel tables.



*Internet Explorer 11 does not support copying and pasting in **Manual entry**.*

There are a number of keyboard shortcuts you can use to work effectively and easily in **Manual entry**.

Shortcuts behavior varies depending if you are selecting cells, rows, or columns, or if you are editing cells in the table. The following table contains the selecting shortcuts:

3 Managing data in the app with Data manager

Keyboard shortcuts for selecting

Shortcut	Description
Arrow keys	Navigates between cell selection
Tab	Moves the cell selection right. If no cell exists to the right, it moves to the first cell in the next row.
Shift+Tab	Moves the cell selection left. If no cell exists to the left, it moves to the first cell in the previous row.
Enter	Toggles to editing mode
Delete	Deletes the current selection

The following table contains the editing shortcuts:

Keyboard shortcuts for editing

Shortcut	Description
Arrow keys	Moves the cursors in the cell.
Tab	Commits the edit and moves to the next cell to the right
Shift+Tab	Commits the edit and moves to the previous cell to the left
Enter	Commits the edit and moves to the next cell below
Shift+Enter	Commits the edit and moves to the previous cell above
Esc	Cancel the edit and toggles to selecting mode

Tables created using **Manual entry** can be edited later to add or remove content. For more information, see *Updating a table from the data source (page 49)*.

Adding data manually

Do the following:

1. Open an app.
2. Open the **Data manager** and then click **+**.
You can also click **Add data** in the **☰** menu.
3. Under **In-App**, click **Manual entry**.
4. Type a name for the table.
5. In the table editor, enter your data.
Double-click a cell to start entering data in the cell.
While editing a cell, clicking any other cell in the table commits your edit and selects the other cell.
6. When your data is complete, click **Add data**.

The table is added to **Data manager**.

Attaching data files and adding the data to the app

You can attach data files to your app, and then use the data in your app.



This functionality is not available in Kubernetes.

An attached file is only available in the app that it is attached to. There is no connection to your original data file, so if you update the original file you need to refresh the attached file.



To avoid exposing restricted data, remove all attached files with section access settings before publishing the app.

Attached files are included when the app is published. If the published app is copied, the attached files are included in the copy. However, if section access restrictions have been applied to the attached data files, the section access settings are not retained when the files are copied, so users of the copied app will be able to see all the data in the attached files.

Limitations

- The maximum size of a file attached to the app is 50 MB.
- The maximum total size of files attached to the app, including image files uploaded to the media library, is 200 MB.
- It is not possible to attach files in Qlik Sense Desktop.

Attaching several data files quickly

The quickest and in most cases the easiest way to attach and add a set of data files to your app is to just drop the files in the app.

Do the following:

- Drop one or more data files in your app.
The files are uploaded and attached to the app, and added to the data model.

When you attach files this way Qlik Sense will try to select the optimal settings for loading the data, for example, recognizing embedded field names, field delimiters or character set. If a table is added with settings that are not optimal you can correct the settings by opening the table in the table editor, and clicking **Select data from source**.



It is not possible to drop files in the data load editor or in the data model viewer.

Attaching a single data file

You can attach data files one by one. This way you gain more control over file import settings, for example, embedded field names, field delimiters or character set used.

3 Managing data in the app with Data manager



Do not add a table in **Data manager** that has already been added as a scripted table with the same name and same columns in **Data load editor**.

Do the following:

1. Open an app.
2. Open the **Data manager** and then click **+**. You can also click **Add data** in the **☰** menu.
3. Drop a data file, or click and select a file from your computer to load.
If you try to attach a file with the same name as an already attached file, you get the option to replace the attached file with the new file.



Each attached file needs to have a unique file name.

4. Select the tables and fields to load.
5. Optionally, select to apply a data filter if you want to select a subset of the data contained in the fields you have selected.

If your data source is a file, select **Filters**. Beside the table to which you want to add a filter, click **Add filter**, select a field, select a condition, and then enter a value with which to filter.



Qlik Sense does not support filters on date fields from QVD files.

Note the following:

- You can apply multiple filters to the same field.
 - You can remove filters in the **Associations** view of **Data manager** or from **Select data from source**. For the changes to take effect, reload data by clicking the **Load data** button.
6. Click **Add data** to open the data in the **Associations** view of the data manager. This allows you to continue to add data sources, transform the data, and associate the tables in **Data manager**. Data profiling is enabled by default when you click **Add data**. Data profiling does the following:
 - Recommends data associations.
 - Auto-qualifies common fields between tables. This adds a unique prefix based on table name.
 - Maps date and time fields to autoCalendar.

Tables are not associated on common field names automatically. You can associate tables in the **Associations** view.



If you want to load the data directly into your app, click **•••** and then disable data profiling. This will also reload all existing data from data sources when you add the data. Tables will be associated on common field names automatically. Date and time fields will not be created.


3 Managing data in the app with Data manager

7. Click **Load data** when you are done preparing the data. If serious problems are detected, you need to resolve the problems in **Data manager** before you can load data into the app.

Deleting an attached file

When you delete a table based on an attached file in the data manager, the table is deleted from the data model, but the attached data file remains in the app. You can delete the data file from the app permanently.

Do the following:

1. Open an app.
2. Open the **Data manager** and then click **+**.
3. Click  **Attached files**.
4. Delete the appropriate file.


The data file is now permanently deleted from the app.



*If you delete an attached file that is used in the app, you will not be able to reload the app until you have removed references to the file in **Data manager**, or in the load script. You edit load scripts in **Data load editor**.*

Reload data from an attached file



A file that you upload for an app is attached to the app. It is only available to that app.

There is no connection to your original data file. If you update the original file, you need to refresh the file that is attached to the app. You can then load the updated data into the app. After reloading the data in **Data manager**, click  (**Refresh data from source**) to see the updated data in the table view.



*Do not add a table in **Data manager** that has already been added as a scripted table with the same name and same columns in **Data load editor**.*

Do the following:

1. Open an app.
2. Open the **Data manager** and then click **+**.
3. Click  **Attached files**.
4. Replace the existing file. The updated file needs to have the same name as the original file. The content of the data file is refreshed.
5. Click **Add data**. Ensure that data profiling is enabled by clicking **...**.
6. In the **Associations** view or the **Tables** view, click the table.
7. Click  to update the data.
8. Click **Load data** to reload the data into the app.



If you have made changes to the field structure of the data file, that is, removed or renamed fields, this can affect the data model in your app, especially if this involves fields that are used to associate tables.

Selecting data fields

You can select which tables and fields to use when you add data, or when you edit a table. Some data sources, such as a CSV file, contain a single table, while other data sources, such as Microsoft Excel spreadsheets or databases, can contain several tables.

If a table contains a header row, field names are usually automatically detected, but you may need to change the **Field names** setting in some cases. You may also need to change other table options, such as **Header size** or **Character set**, to interpret the data correctly. Table options are different for different types of data sources.

Selecting data from a database

The steps for selecting data from a database depend on how you connect to the database. You can connect to an ODBC driver as a DSN source, or you can connect directly through a Qlik Database connector that is part of the Qlik ODBC Connector Package installed with Qlik Sense.

For more information, see *ODBC (page 190)*.

When you add data from a database, the data source can contain several tables.

Do the following:

1. Select a **Database** from the drop-down list.
Some selection dialogs do not have a **Database** drop-down list because the database name is entered when the connection is configured.
2. Select **Owner** of the database.
The list of **Tables** is populated with views and tables available in the selected database. Some databases do not require that owners be specified in the data selection process.
3. Select a table.
4. Select the fields you want to load by checking the box next to each field you want to load.
You can select all fields in the table by checking the box next to the table name.



You can edit the field name by clicking on the existing field name and typing a new name. This may affect how the table is linked to other tables, as they are joined on common fields by default.

5. Select additional tables if you want to add data from them.
6. When you are done with your data selection, click **Add data** to continue with data profiling, and to see recommendations for table relationships.

3 Managing data in the app with Data manager

If you want to load the data directly into your app, click **•••** beside **Add data** and then disable data profiling. This will load the selected data as it is, bypassing the data profiling step, and you can start creating visualizations. Tables will be linked using natural associations, that is, by commonly-named fields.

Selecting data from a Microsoft Excel spreadsheet

When you add data from a Microsoft Excel spreadsheet, the file can contain several sheets. Each sheet is loaded as a separate table. An exception is if the sheet has the same field/column structure as another sheet or loaded table, in which case the tables are concatenated.

Do the following:

1. Make sure you have the appropriate settings for the sheet:

Settings to assist you with interpreting the table data correctly

UI item	Description
Field names	Set to specify if the table contains Embedded field names or No field names . Typically in an Excel spreadsheet, the first row contains the embedded field names. If you select No field names , fields will be named A,B,C...
Header size	Set to the number of rows to omit as table header, typically rows that contain general information that is not in a columnar format.

Example

My spreadsheet looks like this:

Spreadsheet

Machine:	AEJ12B	-	-
Date:	2015-10-05 09	-	-
Timestamp	Order	Operator	Yield
2015-10-05 09:22	00122344	A	52
2015-10-05 10:31	00153534	A	67
2015-10-05 13:46	00747899	B	86

In this case you probably want to ignore the two first lines, and load a table with the fields **Timestamp**, **Order**, **Operator**, and **Yield**. To achieve this, use these settings:

3 Managing data in the app with Data manager

Settings to ignore the two first lines and load the fields

UI item	Description
Header size	2 This means that the first two lines are considered header data and ignored when loading the file. In this case, the two lines starting with Machine: and Date: are ignored, as they are not part of the table data.
Field names	Embedded field names. This means that the first line that is read is used as field names for the respective columns. In this case, the first line that will be read is the third line because the two first lines are header data.

2. Select the first sheet to select data from. You can select all fields in a sheet by checking the box next to the sheet name.
3. Select the fields you want to load by checking the box next to each field you want to load.



You can edit the field name by clicking on the existing field name and typing a new name. This may affect how the table is linked to other tables, as they are joined by common fields by default.

4. When you are done with your data selection, click **Add data** to continue with data profiling, and to see recommendations for table relationships.

If you want to load the data directly into your app, click **•••** beside **Add data** and then disable data profiling. This will load the selected data as it is, bypassing the data profiling step, and you can start creating visualizations. Tables will be linked using natural associations, that is, by commonly-named fields.

Selecting data from a table file

You can add data from a large number of data files.

Do the following:

1. Make sure that the appropriate file type is selected in **File format**.
2. Make sure you have the appropriate settings for the file. File settings are different for different file types.
3. Select the fields you want to load by checking the box next to each field you want to load. You can also select all fields in a file by checking the box next to the sheet name.



You can edit the field name by clicking on the existing field name and typing a new name. This may affect how the table is linked to other tables, as they are joined by common fields by default.

4. When you are done with your data selection, click **Add data** to continue with data profiling, and to see recommendations for table relationships.



If you want to load the data directly into your app, click **•••** beside **Add data** and then disable data profiling. This will load the selected data as it is, bypassing the data profiling step, and you can start creating visualizations. Tables will be linked using natural associations, that is, by commonly-named fields.

Choosing settings for file types

Delimited table files

These settings are validated for delimited table files, containing a single table where each record is separated by a line feed, and each field is separated with a delimited character, for example a CSV file.

File format settings

File format settings for delimited table files

UI item	Description
File format for delimited table files	Set to Delimited or Fixed record . When you make a selection, the select data dialog will adapt to the file format you selected.
Field names	Set to specify if the table contains Embedded field names or No field names .
Delimiter	Set the Delimiter character used in your table file.
Quoting	Set to specify how to handle quotes: None = quote characters are not accepted Standard = standard quoting (quotes can be used as first and last characters of a field value) MSQ = modern-style quoting (allowing multi-line content in fields)
Header size	Set the number of lines to omit as table header.
Character set	Set character set used in the table file.
Comment	Data files can contain comments between records, denoted by starting a line with one or more special characters, for example //. Specify one or more characters to denote a comment line. Qlik Sense does not load lines starting with the character(s) specified here.
Ignore EOF	Select Ignore EOF if your data contains end-of-file characters as part of the field value.

Fixed record data files

Fixed record data files contain a single table in which each record (row of data) contains a number of columns with a fixed field size, usually padded with spaces or tab characters.

Setting field break positions

You can set the field break positions in two different ways:

- Manually, enter the field break positions separated by commas in **Field break positions**. Each position marks the start of a field.

Example: 1,12,24

- Enable **Field breaks** to edit field break positions interactively in the field data preview. **Field break positions** is updated with the selected positions. You can:
 - Click in the field data preview to insert a field break.
 - Click on a field break to delete it.
 - Drag a field break to move it.

File format settings

File format settings for fixed record data files

UI item	Description
Field names	Set to specify if the table contains Embedded field names or No field names .
Header size	Set Header size to the number of lines to omit as table header.
Character set	Set to the character set used in the table file.
Tab size	Set to the number of spaces that one tab character represents in the table file.
Record line size	Set to the number of lines that one record spans in the table file. Default is 1.

HTML files

HTML files can contain several tables. Qlik Sense interprets all elements with a <TABLE> tag as a table.

File format settings

File format settings for HTML files

UI item	Description
Field names	Set to specify if the table contains Embedded field names or No field names .
Character set	Set the character set used in the table file.

XML files

You can load data that is stored in XML format.

There are no specific file format settings for XML files.

QVD files

You can load data that is stored in QVD format. QVD is a native Qlik format and can only be written to and read by Qlik Sense or QlikView. The file format is optimized for speed when reading data from a Qlik Sense script but it is still very compact.

There are no specific file format settings for QVD files.

QVX files

You can load data that is stored in Qlik data eXchange (QVX) format. QVX files are created by custom connectors developed with the Qlik QVX SDK.

There are no specific file format settings for QVX files.

KML files

You can load map files that are stored in KML format, to use in map visualizations.

There are no specific file format settings for KML files.

Returning to the previous step (**Add data**)

You can return to the previous step when adding data.

Do the following:

- Click the back arrow to return to the previous step of **Add data**.

Filtering data from files

You can create filter conditions when you load data from files. This allows you to select a subset of data to load. This can be useful when you want to reduce the amount of data loaded, or only use specific data, such as only sales over \$40,000.

The first time that you add data from a file in the **Add data** step, you can apply filter conditions by clicking **Filters**.

Subsequently, you can change the conditions by clicking your table in the **Data manager**, and then clicking **Edit this table**. Click **Select data from source**, and then click **Filters**.

The available filter conditions are:

- =
- >
- <
- >=
- <=

Consider the following when filtering data. Examples are provided below.

3 Managing data in the app with Data manager

- You can apply filter conditions to numbers, dates, or text.
- Wildcard characters are not supported.
- You can apply multiple conditions. However, conflicting conditions on the same field may result in no data being returned.
- Conditions are applied alphabetically to text data. Conditions are case sensitive.
- You can use more than one letter for text data. For example, `>ct` will return the word *cup*, as will `>=cu`. Note that `>c` will also return *cup*.
- When you use more than one `=` condition, all must evaluate to true to return values. However, when you use more than one `=` condition on the same field, all values that evaluate to true are returned.
- The `<` and `>` conditions, when combined, must all evaluate to true to return values. If these conditions are combined with `=`, all conditions must evaluate to true.
- The `<=` and `>=` conditions, when combined, must all evaluate to true to return values. If these conditions are combined with `=`, all conditions must evaluate to true.
- Filters on date fields from QVD files are not supported.

Examples

These examples use the following values from a single field (one column in a table): *cup*, *fork*, and *knife*.

- Conditions:
 - `=cup`
 - `=fork`
 - `=knife`
 - Returns: *cup*, *fork*, *knife*
 - The equals condition returns all values that are true.
- Conditions:
 - `>b`
 - `<d`
 - Returns: *cup*
 - The letter *c* is both greater than *b* and lesser than *d*.
- Conditions:
 - `<b`
 - `>d`
 - Returns: no values
 - There can be no values that are both lesser than *b* and greater than *d*.
- Conditions:
 - `=fork`
 - `>g`
 - Returns: no values
 - There can be no values that are both equal to *fork* and greater than *g*.

Filtering data from data connectors



Filter data is not currently available for all connectors or Qlik DataMarket.

You enter a data filter expression by selecting **Filter data** in the **Select data to load** step. Selecting **Filter data** opens a text box where you enter a filter expression. For example:

Sales >= 40000

Filter data selects from individual fields, such as *Sales*. It operates as an SQL WHERE clause. Most operators and keywords used in WHERE clauses can be used with **Filter data**. Valid operators include:

- =
- >
- >=
- <
- <=
- IN
- BETWEEN
- LIKE
- IS NULL
- IS NOT NULL

Qlik Sense builds a WHERE clause in the data load script from the expression entered in **Filter data**.

The AND operator can be used to combine operators, such as when you want to filter across more than one field. For example:

Sales <= 30000 AND RegionID = 45

The OR operator can be used to filter data that matches either condition. For example:

Name = 'Smith' OR Name = 'Jones'

You can get the same results from the IN operator. The IN operator is a shorthand method for using multiple OR conditions. For example:

Name IN ('Smith', 'Jones')

Adding data from Qlik DataMarket

Qlik DataMarket offers an extensive collection of up-to-date and ready-to-use data from external sources accessible directly from within Qlik Sense. Qlik DataMarket provides current and historical weather and demographic data, currency exchange rates, as well as business, economic and societal data.



This functionality is not available in Kubernetes.

3 Managing data in the app with Data manager

Qlik DataMarket also offers data sets from the Eurostat database, including Database by themes, Tables by themes, Tables on EU policy, and Cross cutting topics.

When adding data from Qlik DataMarket, you select categories and then filter the fields of data available in those categories. The DataMarket categories contain large amounts of data, and filtering allows you to take subsets of the data and reduce the amount of data loaded.

Some Qlik DataMarket data is available for free. Data packages marked **Premium** are available for a subscription fee.

Before you can use Qlik DataMarket data, you must accept the terms and conditions for its use. Also, if you have purchased a license for premium data packages, you must enter your access credentials to use data in those packages. Once access credentials have been applied, the premium data is labeled **Licensed**.

If you accept the terms and conditions but do not enter a license for any of the premium data packages, the premium packages have a **Purchase** button next to them that enables you to buy a license. The **Purchase** button replaces the **Premium** label.



It is not necessary to accept Qlik DataMarket terms and conditions when using Qlik Sense Desktop. Access credentials are also not required because the premium data sets are not available on Qlik Sense Desktop.

Do the following:

1. Open your app.
2. Open the **Data manager** and then click **+**. You can also click **Add data** in the **☰** menu.
3. Click **Qlik DataMarket** to display the Qlik DataMarket packages.



The DataMarket user interface can be blocked by browser extensions, such as Privacy Badger, that block ads and enhance privacy. This occurs if the extension mistakes DataMarket's communications for user-tracking by a third party. If you encounter this, you can access DataMarket by excluding your Qlik Sense site from the list of blocked sites in the browser extension that blocks DataMarket.

4. Select a data package.
You select a data package from the list (for example, **Essentials Free** or **Financial Reports**), or you can search for specific types of data. To search Qlik DataMarket, select **Search DataMarket** at the top of the list of data packages and enter a search term or phrase in the search text box on the **Select a dataset** step.
5. Select a data set in the **Select a dataset** step.
Depending on the DataMarket package you select, you might have to select several subcategories within the package to get to individual data sets that contain data. Data sets are distinguished by the **i** icon next to the data set description. Clicking the **i** icon displays the data set's metadata. Selecting the data set creates a connection to that specific data collection.

3 Managing data in the app with Data manager

6. Select at least one filter from each dimension, measure and time period in the **Select data to load** step. The left pane lists the dimensions, measures and time periods. If you click on a dimension, measure or time period in the left pane, the values of that dimension, measure or time period are displayed in the right pane.

There is a load size indicator at the bottom of the left column that shows approximately how many cells will be loaded with the currently selected fields. The indicator is green when the number is small, and it becomes yellow when the number increases to a size that might noticeably affect the load time. The indicator becomes red when the amount of data is so large that it might not load successfully.

7. Click **Add data** to open the data in the **Associations** view of the data manager. This allows you to continue to add data sources, transform the data, and associate the tables in **Data manager**.

Data profiling is enabled by default when you click **Add data**. Data profiling does the following:

- Recommends data associations.
- Auto-qualifies common fields between tables. This adds a unique prefix based on table name.
- Maps date and time fields to autoCalendar.

Tables are not associated on common field names automatically. You can associate tables in the **Associations** view.



*If you want to load the data directly into your app, click **•••** and then disable data profiling. This will also reload all existing data from data sources when you add the data. Tables will be associated on common field names automatically. Date and time fields will not be created.*

8. Click **Load data** when you are done preparing the data. If serious problems are detected, you need to resolve the problems in **Data manager** before you can load data into the app.

Selecting Qlik DataMarket dimensions, measures and time periods

Qlik DataMarket data sets contain dimensions, measures and time periods. Dimensions define the context of the data, such as country, age group, or gender. Measures contain numeric data that relates to the context in the dimensions. Time periods specify the years or other periods for which to display the measures.



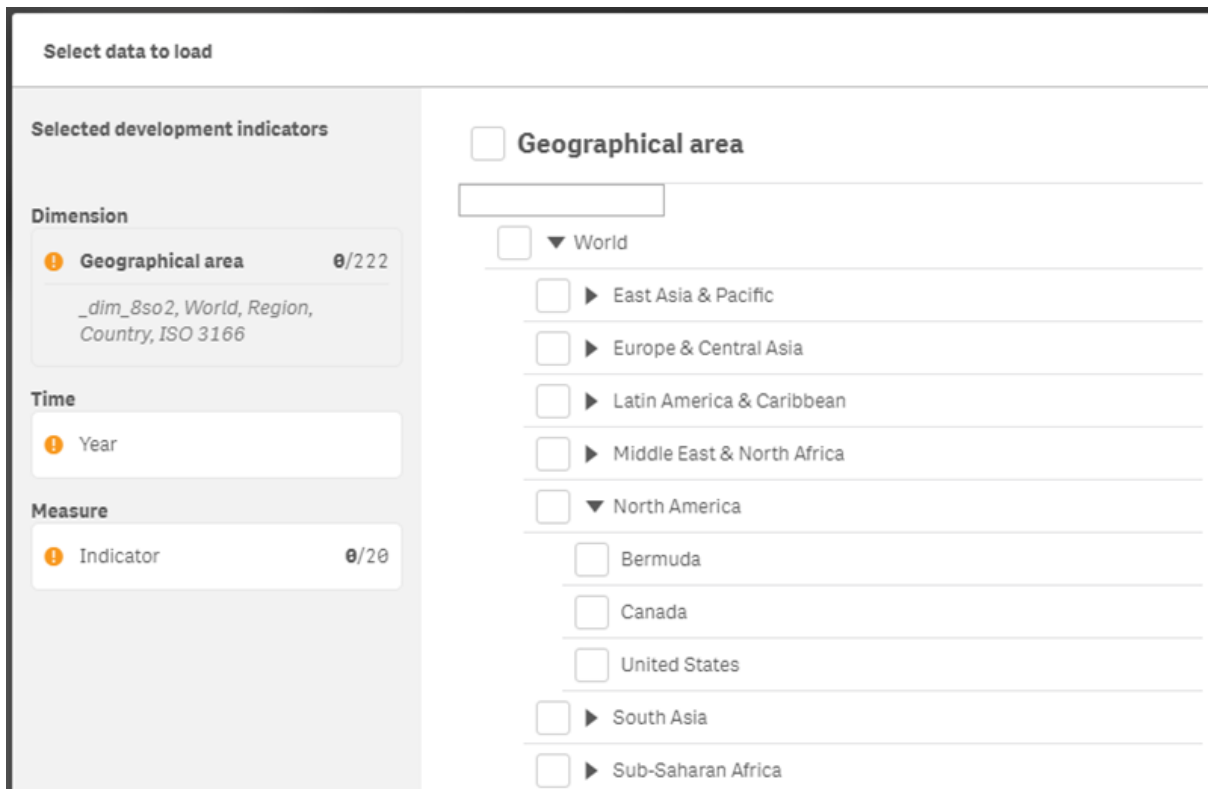
This functionality is not available in Kubernetes.

Data sets contain at least one dimension and one measure, and they all have time dimensions. Before you can add data to an app, you must select at least one dimension and one measure and set the time period. When selecting dimensions, you must include dimensions that contain data. When data is structured hierarchically, it is possible that a parent branch does not contain data.

Some dimensions contain multiple representations of the data. For example, geographical locations designated by country name also contain ISO (International Standards Organization) codes for the countries. Currencies contain regular names, such as Pound sterling and Euro, as well as their ISO 4217 code--GBP and EUR. The extra values for the dimensions are not separately selectable. They are displayed in the description of the dimension.

Selected data to load view with extra values displayed in the description of the dimension.

3 Managing data in the app with Data manager



The extra values are loaded as separate fields.

Geographical area.qlik_datam ...	World	Region	Country	ISO 3166
3	World	Europe & Central Asia		
5	World	East Asia & Pacific		
8	World	North America		
16	World	North America	Canada	CA
41	World	East Asia & Pacific	New Zealand	NZ
k	World	Europe & Central Asia	Austria	AT

In some data sets, it is not necessary to select a measure because the data set contains only one measure. Measure selections are displayed only when there is more than one measure to choose from. For example, the data set **US per capita personal income by state** displays only the geographic dimension and the time period because there is only one measure in the data set--per capita personal income.

There are also data sets that do not require dimension selections. For example, the data sets **US federal interest rate** and **US consumer price index for urban consumers** require only that you select the time period because there is only one dimension and one measure in those data sets. In the first case, the measure is the federal interest rate, and the dimension is the United States. In the second case, the measure is the consumer price index, and the dimension is United States urban consumers.

Data sets may contain only dimensions that have no accompanying data for measures. A data set may contain, for example, only a list of company chief executives (CEOs). In such cases, the dimension is preselected because there are no selections to be made within the dimension.

Selecting hierarchical data from Qlik DataMarket

Many Qlik DataMarket data sets contain dimensions and measures that are structured hierarchically. DataMarket data sets that are structured hierarchically contain two-level and three-level hierarchies. How selections are made in those hierarchies depends on the data at each level.



This functionality is not available in Kubernetes.

Selecting parent and child branches

In some hierarchies, selections can be made at any and all levels. For example, the data set **Selected development indicators** contains the dimension **Geographical area** with three levels.

Selected development indicators contains the dimension Geographical area with three levels.

The screenshot shows the 'Select data to load' interface. On the left, under 'Selected development indicators', the 'Dimension' section shows 'Geographical area' (0/222) with a list of values: '_dim_8so2, World, Region, Country, ISO 3166'. The 'Time' section shows 'Year' (0/20) and the 'Measure' section shows 'Indicator' (0/20). On the right, the 'Geographical area' dimension is expanded, showing a search box and a list of hierarchical levels with checkboxes: 'World', 'East Asia & Pacific', 'Europe & Central Asia', 'Latin America & Caribbean', 'Middle East & North Africa', 'North America' (expanded to show 'Bermuda', 'Canada', 'United States'), 'South Asia', and 'Sub-Saharan Africa'.

A selection from either **World**, **Region**, or **Country** is valid by itself. Any selection that includes the highest level loads all the data for the regions and countries even if specific regions and countries are also selected. But if a region is selected by itself, then only that region of the world is loaded.

If you select both **World** and **North America**, world data is displayed separately from North America data.

If you select **Canada** from **Country**, then you get separate data for the world, the North America region, and Canada.

3 Managing data in the app with Data manager

If you select **Canada** from **Country** but do not select **North America**, then the aggregate data for North America is not loaded. Only the data for Canada is loaded for the North America region.

Selecting parent branches that contain no data

In some data sets, the parent branch in a hierarchy contains no data. Parent branches with no data are indicated by check boxes with broken lines, as in the data set **US social characteristics (by state)**.

US Social characteristics (by state) contains three parent branches with no data; Ancestry, Disability status of the civilian noninstitutionalized population and Educational attainment.

The screenshot shows the 'Select data to load' interface for the data set 'US social characteristics (by state)'. On the left sidebar, under 'Dimension', 'Geography' is selected (1/53). Under 'Time', 'Year' is selected (Most recent). Under 'Measure', 'Subject' is selected (0/164). The main area shows a list of subject categories with checkboxes. The 'Subject' category is selected. The following categories are listed with broken checkboxes, indicating no data: Ancestry, Disability status of the civilian noninstitutionalized population, Educational attainment (with a dropdown arrow), Population 25 years and over, Population 25 years and over - 9th to 12th grade, no diploma, Population 25 years and over - Associate's degree, Population 25 years and over - Bachelor's degree, Population 25 years and over - Graduate or professional degree, Population 25 years and over - High school graduate (includes e...), Population 25 years and over - Less than 9th grade, Population 25 years and over - Some college, no degree, and Fertility.

When the parent field is selected, all the children of the branch are automatically selected as well. To select some rather than all of the children in the branch, you can either deselect individual fields from the automatic parent selection or select individual fields without selecting the parent field.

Selecting aggregate data from Qlik DataMarket

Some Qlik DataMarket data sets contain dimensions that include an aggregate field value. These dimensions represent hierarchies and contain more than one field, representing levels of that hierarchy. Fields within the hierarchy may contain blank values. The numbers associated with those blank values are aggregates of the numbers associated with the other values in that field.

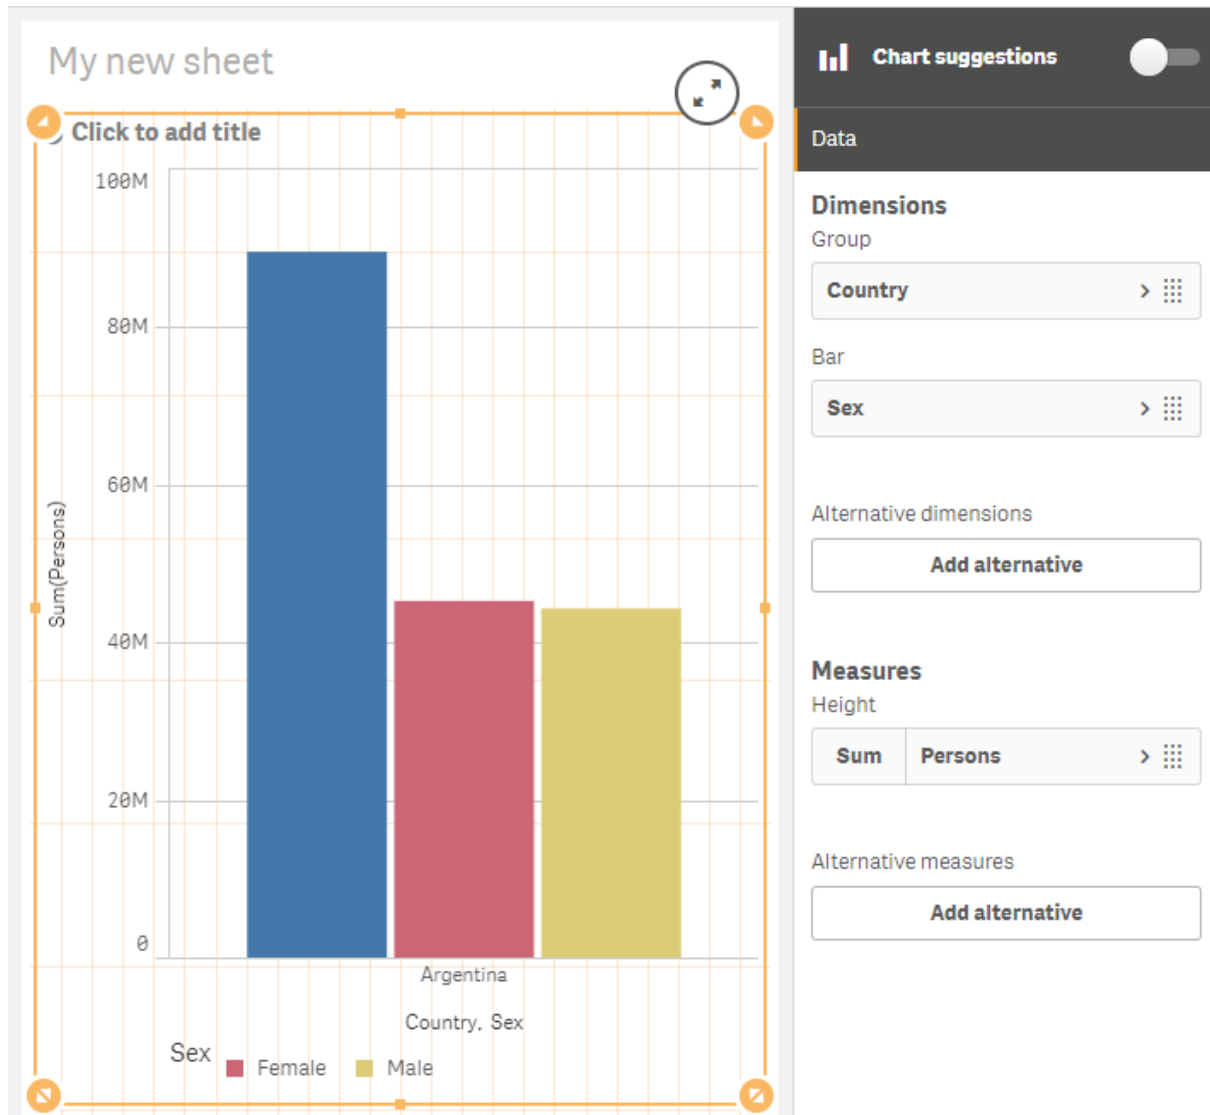


This functionality is not available in Kubernetes.

3 Managing data in the app with Data manager

For example, the **World population by country** data set contains a **Sex** dimension. It has the subset field values of **Female** and **Male**. When **Sex** is loaded into Qlik Sense, it contains three values: **Female**, **Male**, and a blank value for the aggregate field value.

Bar chart with the dimension Sex, the total male and female population for Argentina is displayed as two separate bars with an unnamed bar for the aggregate total.



In the example displayed in the image, the aggregate field value contains the total of both **Female** and **Male**.

The blank value field is included whenever all subset field values of a dimension are included. The totals associated to the aggregate field values are present in the data and can result in double-counting if they are included when calculating aggregates over data.

Depending on your visualizations used, you can exclude the aggregate field values. For example:

- Set the blank value to null using the **Set nulls** data profiling card in **Data manager** and exclude the null value from the dimension by clearing **Include null values** in the **Dimension** section of **Properties**.

3 Managing data in the app with Data manager

- Use an expression to limit what dimension values are included and then clear **Include null values** in the **Dimension** section of **Properties**. For example, from the **Sex** dimension, you could use the expression `=if(match(Sex,'Female','Male'),Sex)` to exclude the aggregate field value.
- Use Set analysis expressions to exclude the blank value's aggregate numbers from a measure.



*When working with data sets such as **World population by country** that contain multiple fields with aggregate data, ensure that the tables with those aggregate fields are not directly associated. If they are directly associated, they are likely to create a circular reference.*

Searching Qlik DataMarket

The search facility in Qlik DataMarket helps you find the data you need by searching on terms or phrases rather than browsing data packages and data sets. The term or phrase entered is searched for in the names and descriptions of DataMarket data sets and in the values within the data sets.



This functionality is not available in Kubernetes.

For example, if you search for the term *europa*, you get a list first of all the data sets with the term *Europe* in the title followed by data sets that contain data labeled with the term. In the case of the term *europa*, one of the data sets found is **Selected development indicators**, which contains the term in its **Geographical area** dimension--*Europe & Central Asia*.

DataMarket searches on the literal term or phrase you enter, and it also searches on related terms or synonyms. A term entered in singular form is also searched for in its plural form. For example, the terms *currency* and *index* have plural forms--*currencies* and *indices*--that are searched for at the same time their singular forms are searched for.

The search facility also looks for matches based on the stem or root form of terms. For example, if you search on the term *production*, the root form of the word--*product*--is also searched for.



*DataMarket does not search partial terms. For example, it does not find the string "prod" even though it is part of the terms *product* and *production*, which are terms the search facility does find in phrases such as *Gross Domestic Product*.*

Qlik DataMarket also contains an index of synonyms, so you can find a wide range of data without using the exact term used in the name or description of the data collection or in the data fields. For example, data sets that use the dimension labeled **Sex** are also found with the term *gender*. The DataMarket search facility has over 200 sets of synonyms. Some synonyms included are:

- earnings, income
- GBP, pound
- health care, healthcare

3 Managing data in the app with Data manager

- labor, labour
- salary, wages, pay, earnings

The search results are displayed from highest relevance to lowest. Relevance is determined by where the search term is found. Terms found in data set names or descriptions rank higher than terms found in data set values.

When multiple search terms are entered, the results do not necessarily include all the terms. If only one of the terms is found, the entry containing the term is returned as one of the search results. However, entries that contain more of the entered search terms rank higher.

Only the top 50 results are listed for a search.

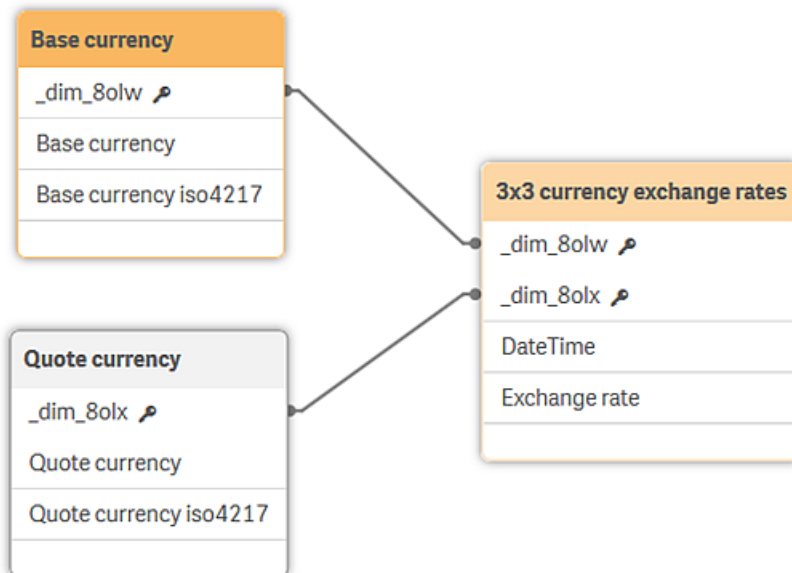
To narrow searches, you can exclude terms from the search by placing a hyphen before the terms when they are entered in the search string. For example, you can search for "US" but exclude unemployment by placing a hyphen before the search term, "-unemployment."

Making associations in Qlik DataMarket



This functionality is not available in Kubernetes.

When data is loaded from a Qlik DataMarket data set, it is allocated to multiple individual tables. These tables are associated by generated key fields. Measures and time periods from the data set are consolidated in one table that is assigned the name of the data set. Dimension fields are allocated to individual tables. For example, the **3x3 currency exchange rates** data set loads as three tables:



Some dimensions offer additional fields when loaded. The extra fields provide additional representations of the dimensions. In the **3x3 currency exchange rates** data set, the currencies are also listed by the iso4217 representation. For example:

3 Managing data in the app with Data manager

Currencies

Base currency	Euro
Base currency iso4217	EUR
Quote currency	US dollar
Quote currency iso4217	USD

Data sets with population data by country and region offer extra dimension representations for the region names, such as ISO 3166 codes.

Associating a data set's multiple tables

All tables from a single data set should be associated through the measures table. When you load a data set from **Add data**, associations are made automatically from the generated key fields. The following data model shows the data set **US population by state, race, sex and age**. The dimension tables (**Age**, **Sex**, **Location**, and **Race**) are associated through the measures table, which has the data set name.



These associations are required to interpret relationships between the dimensions and the measures that are important in visualizations. For example, if a company wants to use the US population data to compare its product sales to age groups in various US states, the **Age** and **Location** dimensions must be associated through the measures table to get the number of people in each group in the various states.

When data sets have multiple dimension tables, there are often additional associations that can be made. For example, aggregate fields usually have the same value ("Total") that suggests a possible association. Such associations are not useful, however, and can result in circular references.

The multiple-table structure increases the efficiency with which data is loaded, and it can improve the data associations.

Converting single-table data sets to multiple tables

Prior to Qlik Sense 3.0, DataMarket data sets were loaded as single tables. In Qlik Sense 3.0, those previously loaded tables remain as single tables, and any visualizations that use them continue to work as before. You can convert those single tables to multiple tables, and the visualizations will continue to work without modification. By converting them, you gain the efficiency of multiple tables.

Do the following:

3 Managing data in the app with Data manager

1. Open **Data manager**.
2. Delete the table containing the DataMarket data set you want to convert to multiple tables.
The table is marked **Pending delete**.
3. Click **+** to add data.
4. Select **Qlik DataMarket**.
5. Select the data set corresponding to the table you deleted.
When the **Select data to load** page opens, the selection you made previously is displayed. Do not change the selections unless you also plan to change the visualizations that use the data set. If you change the selections, existing visualizations that use the data set may no longer work.
6. Click **•••** beside the **Add data** button and then disable data profiling.
7. Click **Add data**.
The data set is loaded in multiple tables and the single table is deleted.

Check the visualizations that use the data set you converted to multiple tables. They should work as they were originally designed unless you changed the data selected either by adding or removing some of the selections made when the table was loaded previously.

Integrating corporate data with DataMarket data

Qlik DataMarket data can be integrated with your corporate data to provide a broader context for that data and yield greater insights into your markets.



This functionality is not available in Kubernetes.

DataMarket data comes from a variety of sources, and as a result, associations with your data may not always be immediately evident. You might find that a number of associations have to be edited in the data preparation step. For example, you might find it valuable to evaluate certain characteristics of the countries you operate in. But fields for countries in some DataMarket data sets might not have enough values in common with your corporate data to make the association useful. That is why you must carefully assess the associations between your data and DataMarket data.

The following illustrations demonstrate how to integrate corporate and DataMarket data and create meaningful Qlik Sense visualizations.

Compare sales to Gross Domestic Product (GDP)

When you are doing business in a country or considering entering a country's market, there are a number of indicators that can help you determine if your company's performance aligns as you expect with the demographic and economic environment in that country.

The corporate data in this illustration enables sales data to be aggregated by country. A bar chart compares the sales by country. That can provide insight into how your company is doing across all its markets.

To see how you are doing within each country, you could compare your company's sales to country data that indicates how strong the market is. For example, you could compare sales in each country with the country's

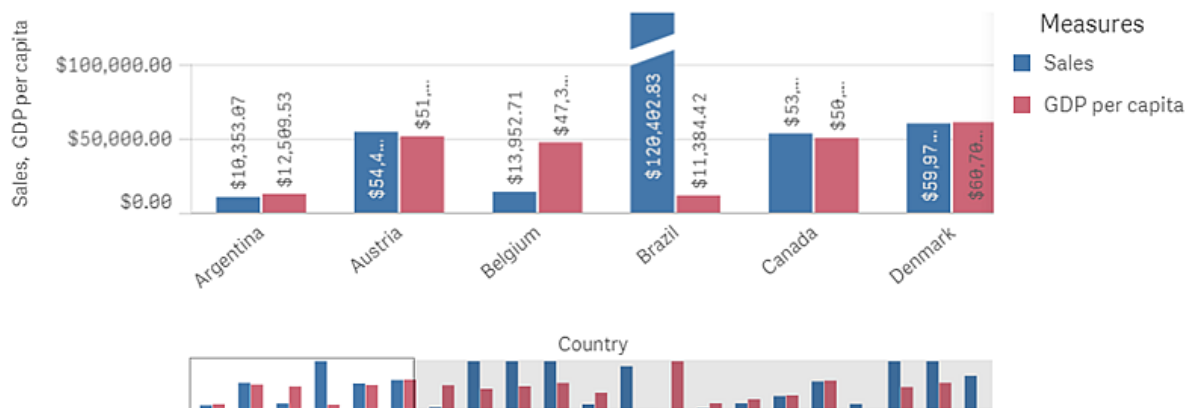
3 Managing data in the app with Data manager

Gross Domestic Product (GDP). Or you could compare sales to the demographics of your target market. If your company's target is people ages 21 to 35 year, you can see how many people the countries have in that age category. Or see what percentage of the total population is in that age category.

Qlik DataMarket contains a data set in the Essentials Free group called *Select development indicators* that provides a number of economic measures, including GDP growth rates, literacy, internet users, total population, and GDP per capita in US dollars. To associate country data from *Select development indicators*, the corporate data must have a field that matches the Country field in the DataMarket data set. If the corporate data has many more countries than the *Select development indicators*, the association would not be strong and probably not useful. If the corporate data has fewer countries than the *Select development indicators*, the association can probably be useful for a comparison.

Assuming there is a good association between the country fields in the corporate and DataMarket data, you can add GDP per capita in US dollars to the sales bar chart to compare the sales in each country to GDP.

Sales per country compared to per capita GDP



Show sales in euros

When reviewing sales data, a common visualization is a KPI chart that shows total sales. In the illustration here, the sales total is in US dollars because that is how the corporate data is recorded. With so many sales in countries that use the euro, it might be good to see the sales total in euros. That can be added to the KPI. In the Essentials Free group, DataMarket contains *3x3 currency exchange rates* that can be used to convert US dollars to euros.

When selecting the Base currency field in the *3x3 currency exchange rates* data set, you would select only US dollars because that is the currency in which sales are recorded. In the corporate data set used in this illustration, there is a field named Base currency that indicates the currency each customer uses. However, it contains twelve different currencies, and as a result, the **Data manager** recommends against associating the two fields. You should not associate those fields as the currencies in the corporate data that do not correspond to US dollars and euros can interfere with some comparisons of dollars and euros. The data model then appears as follows:

3 Managing data in the app with Data manager



The Quote currency from *3x3 currency exchange rates* should be Euro. The DateTime selection should be Most recent because you want only the current exchange rate, not historical data, for the KPI visualization.




To get Euro Sales, you simply multiply the `Sum(Sales)` by the `ExchangeID` rate in the *3x3 currency exchange rates* data set.

Create alternative measures for sales

With DataMarket currency data, you can add a further enhancement to the sales bar chart. The chart compares sales to per capita GDP in US dollars for each country, but alternative measures can be added for the sales and GDP. The `Sum(Sales)` can be multiplied by the `Exchange rate` in the *3x3 currency exchange rates* table to show sales in euros, just as was done in the Sales KPI. The per capita GDP in US dollars can be multiplied by the `Exchange rate` as well, and then the bar chart compares sales to per capita GDP in euros.

Sales per country compared to per capita GDP





3 Managing data in the app with Data manager

This is where it is important that the Base currency fields not be linked because, as noted above, the corporate data set's Base currency field contains twelve different currencies. When the exchange-rate calculation is performed on the separate countries, the base currency for each country would be used if the tables were linked. But the corporate data does not contain any sales values in most of those twelve currencies. It contains sales values only in US dollars. And the Base currency from the DataMarket data set is only US dollars, so for any country that has a Base currency value other than US dollars in the corporate data, Sales in euros would be a null value if the two tables were linked.

3.12 Editing a table

You can edit a table that was added to the app in the **Data manager** overview, to rename the table, associate the table to other tables, or make field transformations.

To edit a table, select the table in **Data manager** and click . The table editor is displayed, with a preview of the data in the table. Each field has a field menu with transformation options. You open the field menu by clicking . Selecting a field displays the data profiling card pane, which contains a summary of the field's data as well as additional transformation options.



If the data contains records with identical data in all fields that are loaded, they are represented by a single record in the preview table.

Renaming a table

When you add a table in **Data manager**, the table is assigned a default name, based on the name of the database table, data file, or Excel worksheet, for example. If the name is non-descriptive or unsuitable, you can rename it.

Do the following:

1. Click on the table name.
2. Edit the table name.
3. Press Enter or click outside the table name.

The table is now renamed.

Renaming a field

You can rename fields in a table to get a better name that is easier to understand.

Do the following:

1. Click on the field name that you want to rename, or select **Rename** from the field menu.
2. Type the new name.



Field names must be unique. If you have fields with the same name in several tables, Qlik Sense will qualify the field names when you add data, that is, add the table name as prefix.

3. Press the Enter key, or click outside the field.

The field is now renamed.

Managing associations to other tables

You can create custom associations to fields in other tables with **Associate** in the field menu.

Typically, these are the most common cases where you need to create a custom association instead of following the recommendations:

- You know which fields to associate the tables with, but the score for this table pair is too low to show in the list of recommendations.
Create an association based on a single field in each table.
- The tables contain more than one common field, and they need to be used to form the association.
Create a compound key.

In many cases it is easier to manage to your associations in the **Associations** view

Changing field type and display format

When data is added, Qlik Sense interprets the field type of each field. The following field types are currently supported:

-  **General**
-  **Date**
-  **Timestamp**
-  **Geo data**

If the data was not interpreted correctly, you can adjust the field type. You can also change the input and display format of a date or timestamp field.

Fields that contain geographical information in the form of names or codes, such as postal areas, cannot be used for mapping unless they are designated as **Geo data** fields.

Hiding fields from analysis

You can hide fields from sheet view or insight advisor so that they are only available in **Data manager** and **Data load editor**. You might, for example, have fields that are only used for calculating another field. You can hide these fields so that they are not available in the assets panel of sheets or from insight advisor but remain available in **Data manager**.

For more information, see *Hiding fields from analysis (page 67)*.

Assessing table field data before loading data

You can examine the data in your table for potential quality issues such as null values and outlier values before you load it using the **Summary** data profiling card. The **Summary** card categorizes fields as dimensions, measures, or temporal fields, providing different data summaries for each and enabling different transformation options in other data profiling cards. Fields set as measures in the **Summary** card can be grouped using the **Bucket** card. Fields set as dimensions in the **Summary** card can have a custom order applied in the **Order** card. For fields that can be classified in multiple categories, you can switch between each possible category's summary for the field.

For more information, see *Assessing table field data before loading data (page 68)*.

Replacing field values in a table

You can replace values in a field using the **Replace** data profiling card. The **Replace** card enables you to select one or more values from a field and then replace them with another value. For example, in a data set that contains country names in both full and abbreviated formats, you could replace them with a single shared value.

For more information, see *Replacing field values in a table (page 69)*.

Setting field values as null in a table

You can set distinct values from a dimension field to be treated as null values using the **Set nulls** data profiling card. The **Set nulls** card enables you to select values from a table field and then manually set them as null. For example, if your data set represents nulls using a character such as X, you can use the **Set nulls** card to enable Qlik Sense to treat that value as null. The **Set nulls** card can also be used to clean your data set by setting unwanted values as nulls.

For more information, see *Setting field values as null in a table (page 71)*.

Setting a custom order for field values

Depending on your data, it may be more meaningful to display dimension values in an order other than alphabetical or numerical. Fields set as dimensions in the **Summary** data profiling card may have a custom order of data applied using the **Order** data profiling card, enabling you to set the default organization of field data in visualizations.

For more information, see *Customizing the order of dimension values (page 72)*.

Splitting a field in a table

You can extract information from an existing field into new fields using the **Split** data profiling card. The **Split** card enables you to split content from a field into multiple fields. You could, for example, split a field that contains an address to extract the zip or postal code. This enables you to quickly create new fields containing relevant sections of existing data.

For more information, see *Splitting a field in a table (page 73)*.

Grouping measure data into ranges

You can group values in a table measure field into ranges using the **Bucket** data profiling card. The **Bucket** card enables you to group a field's value in user-defined buckets, creating a new field that is added to the table. You could, for example, group ages into age ranges to use as a dimensions in your visualizations.

Viewing field transformation details

You can view the current operations and transformations performed on a field and their sequence in the **Details** dialog. **Details** enables you to understand where a field came from, what changes have been made to it, and the sequence in which the transformations were applied.

For more information, see *Viewing table and field transformation details in Data manager (page 84)*.

Unpivoting crosstab data

If you have loaded data in crosstab format, the best option is usually to unpivot the table, that is, transposing parts of the table into rows. This will make it easier to work with the data and create associations to your other data tables.

Year	Europe	RoW
2007	234	567
2008	345	534

→

Year	Region	Sales
2007	Europe	234
2007	RoW	567
2008	Europe	345
2008	RoW	534

For more information, see *Unpivoting crosstab data in the data manager (page 78)*.

Updating a table from the data source

You may want to change the selection of fields from the data source. For example, you may need to add a field that was left out, or the data source may have been updated with added fields. In this case, you can update the table from the data source. If the table was created with **Manual entry**, you can add, edit, or delete table data as well as add new rows and columns. For more information, see *Adding data manually in Qlik Sense (page 21)*.

Do the following:

1. Click **Select data from source**.
The data selection wizard opens with your current selections.
2. Make the required changes in selection.
3. Click **Add data** with data profiling enabled.

The table is now updated with fields according to the selections you made.

Adding a calculated field

There are many cases where you need to adjust or transform the field data that is loaded. For example, you may need to concatenate a first name and a last name to a full name, extract part of a product number, convert the data format or multiply two numbers.

You can add calculated fields to manage many cases like this. A calculated field uses an expression to define the result of the field. You can use functions, fields and operators in the expression. You can only refer to fields in the table that you are editing.

Sorting a table

You can sort a table based on a specific field while you are editing the table, to get a better overview of the data. You can only sort on one field at a time.

Do the following:



- From the field menu, select **Sort**.

The table data is now sorted in ascending order according to this field. If you want to sort in descending order, select **Sort** again.



The sort order is not maintained in the loaded app data.

Undo and redo actions

You can undo and redo your table edit actions by clicking  and .

The undo/redo history is cleared when you close the table editor.

Associating data in the table editor

You can create custom associations to fields in other tables with **Associate** in the field menu of the **Data manager** table editor.

In many cases it is easier to manage your associations in the **Associations** view.




Typically, these are the most common cases where you need to create a custom association instead of following the recommendations:



- You know which fields to associate the tables with, but the score for this table pair is too low to show in the list of recommendations.
Create an association based on a single field in each table.
- The tables contain more than one common field, and they need to be used to form the association.
Create a compound key.

Creating an association using a single field

If the two tables contain related data, but the association does not show up as recommended, you can define a custom association in the table editor. This creates a key field to associate the tables.

Do the following:




1. From the data manager overview, click  on one of the tables you want to associate.
The table editor opens.
2. Select **Associate** from the field menu of the field you want to use in the key field.
The **Associate tables** editor opens, with a preview of the field you selected in the left table. Now you need to select which field to associate this with in the right hand table.
3. Click **Select table** and select the table to associate with.
4. Click  and select the field to associate with.
The right hand table will show preview data of the field you selected. Now you can compare the left table with the right to check that they contain matching data. You can search in the tables with  to compare them more easily.
5. Enter a name for the key field that will be created in **Name**.
It's not possible to use the same name as an existing field in either of the tables.
6. Click **Associate**.

The tables are now associated by the two fields you selected, using a key field. This is indicated with . Click  to display options to edit or break the association.


Creating a compound key

If two tables contain more than one common field that would create an association, Qlik Sense creates a synthetic key to handle the association. The recommended way of fixing this is to create a compound key. This can be achieved by creating a custom association containing all fields that should be associated.

Do the following:

1. From the data manager overview, click  on one of the tables you want to associate.
The table editor opens.
2. Select **Associate** from the field menu of one of the fields you want to include in the compound key field.
The **Associate tables** editor opens, with a preview of the field you selected in the left table.
3. Click  to add the other fields you want to include in the compound key field.
The preview is updated with the compound key data.
Now you need to select which fields to associate this with in the right hand table.
4. Click **Select table** and select the fields you want to include in the compound key field.
5. Click  and select the field to associate with. You need to select them in the same order as in the left hand table.
To make it easier to interpret the data in the key you can also add delimiter characters.
The right hand table will show preview data of the field you selected.

3 Managing data in the app with Data manager

Now you can compare the left table with the right to check that they contain matching data. You can search in the tables with  to compare them more easily.

6. Enter a name for the key field that will be created in **Name**.
7. Click **Associate**.

The tables are now associated by the fields you selected, using a compound key field.

Limitations



There are some limitations to the use of compound keys.

- It is not possible to create a compound key in a concatenated table.
- If you use a calculated field in a compound key, the calculated field expression is expanded in the compound key expression. There is no reference to the calculated field, that is, if you edit the calculated field, the compound key is not updated.

Editing an association

You can edit an association to rename it, or change the associated fields..

Do the following:



1. Click  to show the association menu.
2. Click  to edit the association.

The **Associate tables** editor opens, and you can rename the association or change the associated fields.

Breaking an association

If you have a created an association between two tables that is not needed, you can break it.

Do the following:

1. Click  to show the association menu.
2. Click  to break the association.

The tables are now not associated anymore.

Using calculated fields



There are many cases where you need to adjust or transform the field data that is loaded. For example, you may need to concatenate a first name and a last name to a full name, extract part of a product number, convert the data format or multiply two numbers.

You can add calculated fields to manage many cases like this. A calculated field uses an expression to define the result of the field. You can use functions, fields and operators in the expression. You can only refer to fields in the table that you are editing. You can reference another calculated field in your calculated field.

You add and edit calculated fields in the table editor of the data manager.

Adding a calculated field

Do the following:

1. Click **Add field** and select **Calculated field**.
The editor for **Add calculated field** opens.
2. Type the name of the calculated field in **Name**.
3. Define the expression of the calculated field in **Expression**. There are two different ways to do this.
 - Use the ***fx*** (**Functions**),  (**Fields**) and  (**Operators**) lists to select and insert items into the expression.
The item you select is inserted at the cursor position in **Expression**.
 - Type the expression for the calculated field in **Expression**.
As you type, you get assistance with suggested functions and fields, as well as function syntax.
Example results of the calculated field is displayed in **Preview**.
4. Click **Create** to create the calculated field and close the calculated field editor.

Editing a calculated field

You can change the name or edit the expression of a calculated field.

Do the following:

1. Select **Edit** from the drop-down menu next to the field name.
The editor for **Update calculated field** opens.
2. Edit the name of the calculated field in **Name** if you want to change it.
3. Edit the expression of the calculated field.
4. Click **Update** to update the calculated field and close the calculated field editor.

Which functions can I use in a calculated field expression?

You can use the functions listed here when you create a calculated field expression. This is a subset of the expressions available in the data load script. The expression cannot result in any aggregation of data from several records, or use inter-record functions to refer to data in other records.

String functions that can be used in a calculated field expression

These functions can be used to modify or extract data in text string format.

String functions

Function	Description
Capitalize	Capitalize() returns the string with all words in initial uppercase letters.
Chr	Chr() returns the Unicode character corresponding to the input integer.

3 Managing data in the app with Data manager

Function	Description
FindOneOf	FindOneOf() searches a string to find the position of the occurrence of any character from a set of provided characters. The position of the first occurrence of any character from the search set is returned unless a third argument (with a value greater than 1) is supplied. If no match is found, 0 is returned.
Index	Index() searches a string to find the starting position of the nth occurrence of a provided substring. An optional third argument provides the value of n, which is 1 if omitted. A negative value searches from the end of the string. The positions in the string are numbered from 1 and up.
KeepChar	KeepChar() returns a string consisting of the first string, 'text', less any of the characters NOT contained in the second string, "keep_chars".
Left	Left() returns a string consisting of the first (left-most) characters of the input string, where the number of characters is determined by the second argument.
Len	Len() returns the length of the input string.
Lower	Lower() converts all the characters in the input string to lower case.
LTrim	LTrim() returns the input string trimmed of any leading spaces.
Mid	Mid() returns the part of the input string starting at the position of the character defined by the second argument, 'start', and returning the number of characters defined by the third argument, 'count'. If 'count' is omitted, the rest of the input string is returned. The first character in the input string is numbered 1.
Ord	Ord() returns the Unicode code point number of the first character of the input string.
PurgeChar	PurgeChar() returns a string consisting of the characters contained in the input string ('text'), excluding any that appear in the second argument ('remove_chars').
Repeat	Repeat() forms a string consisting of the input string repeated the number of times defined by the second argument.
Replace	Replace() returns a string after replacing all occurrences of a given substring within the input string with another substring. The function is non-recursive and works from left to right.
Right	Right() returns a string consisting of the last (right-most) characters of the input string, where the number of characters is determined by the second argument.
RTrim	RTrim() returns the input string trimmed of any trailing spaces.
SubStringCount	SubStringCount() returns the number of occurrences of the specified substring in the input string text. If there is no match, 0 is returned.
TextBetween	TextBetween() returns the text in the input string that occurs between the characters specified as delimiters.
Trim	Trim() returns the input string trimmed of any leading and trailing spaces.
Upper	Upper() converts all the characters in the input string to upper case for all text characters in the expression. Numbers and symbols are ignored.

3 Managing data in the app with Data manager

Date functions that can be used in a calculated field expression

Qlik Sense date and time functions are used to transform and convert date and time values.

Functions are based on a date-time serial number that equals the number of days since December 30, 1899. The integer value represents the day and the fractional value represents the time of the day.

Qlik Sense uses the numerical value of the argument, so a number is valid as a argument also when it is not formatted as a date or a time. If the argument does not correspond to a numerical value, for example, because it is a string, then Qlik Sense attempts to interpret the string according to the date and time environment variables.

If the date format used in the argument does not correspond to the one set in the **DateFormat** system variable, Qlik Sense will not be able to interpret the date correctly. To resolve this, either change the settings or use an interpretation function.

Date functions

Function	Description
addmonths	This function returns the date occurring n months after startdate or, if n is negative, the date occurring n months before startdate .
addyears	This function returns the date occurring n years after startdate or, if n is negative, the date occurring n years before startdate .
age	The age function returns the age at the time of timestamp (in completed years) of somebody born on date_of_birth .
converttolocaltime	Converts a UTC or GMT timestamp to local time as a dual value. The place can be any of a number of cities, places and time zones around the world.
day	This function returns an integer representing the day when the fraction of the expression is interpreted as a date according to the standard number interpretation.
dayend	This function returns a value corresponding to a timestamp of the final millisecond of the day contained in time . The default output format will be the TimestampFormat set in the script.
daylightsaving	Converts a UTC or GMT timestamp to local time as a dual value. The place can be any of a number of cities, places and time zones around the world.
dayname	This function returns a value showing the date with an underlying numeric value corresponding to a timestamp of the first millisecond of the day containing time .
daynumberofquarter	Converts a UTC or GMT timestamp to local time as a dual value. The place can be any of a number of cities, places and time zones around the world.
daynumberofyear	This function calculates the day number of the year in which a timestamp falls. The calculation is made from the first millisecond of the first day of the year, but the first month can be offset.

3 Managing data in the app with Data manager

Function	Description
daystart	This function returns a value corresponding to a timestamp with the first millisecond of the day contained in the time argument. The default output format will be the TimestampFormat set in the script.
firstworkdate	The firstworkdate function returns the latest starting date to achieve no_of_workdays (Monday-Friday) ending no later than end_date taking into account any optionally listed holidays. end_date and holiday should be valid dates or timestamps.
GMT	This function returns the current Greenwich Mean Time, as derived from the system clock and Windows time settings.
hour	This function returns an integer representing the hour when the fraction of the expression is interpreted as a time according to the standard number interpretation.
inday	This function returns True if timestamp lies inside the day containing base_timestamp .
indaytotime	This function returns True if timestamp lies inside the part of day containing base_timestamp up until and including the exact millisecond of base_timestamp .
inlunarweek	This function finds if timestamp lies inside the lunar week containing base_date . Lunar weeks in Qlik Sense are defined by counting 1 January as the first day of the week.
inlunarweektodate	This function finds if timestamp lies inside the part of the lunar week up to and including the last millisecond of base_date . Lunar weeks in Qlik Sense are defined by counting 1 January as the first day of the week.
inmonth	This function returns True if timestamp lies inside the month containing base_date .
inmonths	This function finds if a timestamp falls within the same month, bi-month, quarter, tertial, or half-year as a base date. It is also possible to find if the timestamp falls within a previous or following time period.
inmonthstodate	This function finds if a timestamp falls within the part a period of the month, bi-month, quarter, tertial, or half-year up to and including the last millisecond of base_date . It is also possible to find if the timestamp falls within a previous or following time period.
inmonthtodate	Returns True if date lies inside the part of month containing basedate up until and including the last millisecond of basedate .
inquarter	This function returns True if timestamp lies inside the quarter containing base_date .

3 Managing data in the app with Data manager

Function	Description
inquartertodate	This function returns True if timestamp lies inside the part of the quarter containing base_date up until and including the last millisecond of base_date .
inweek	This function returns True if timestamp lies inside the week containing base_date .
inweektodate	This function returns True if timestamp lies inside the part of week containing base_date up until and including the last millisecond of base_date .
inyear	This function returns True if timestamp lies inside the year containing base_date .
inyeartodate	This function returns True if timestamp lies inside the part of year containing base_date up until and including the last millisecond of base_date .
lastworkdate	The lastworkdate function returns the earliest ending date to achieve no_of_workdays (Monday-Friday) if starting at start_date taking into account any optionally listed holiday . start_date and holiday should be valid dates or timestamps.
localtime	This function returns a timestamp of the current time from the system clock for a specified time zone.
lunarweekend	This function returns a value corresponding to a timestamp of the last millisecond of the lunar week containing date . Lunar weeks in Qlik Sense are defined by counting 1 January as the first day of the week.
lunarweekname	This function returns a display value showing the year and lunar week number corresponding to a timestamp of the first millisecond of the first day of the lunar week containing date . Lunar weeks in Qlik Sense are defined by counting 1 January as the first day of the week.
lunarweekstart	This function returns a value corresponding to a timestamp of the first millisecond of the lunar week containing date . Lunar weeks in Qlik Sense are defined by counting 1 January as the first day of the week.
makedate	This function returns a date calculated from the year YYYY , the month MM and the day DD .
maketime	This function returns a time calculated from the hour hh , the minute mm , and the second ss .
makeweekdate	This function returns a date calculated from the year YYYY , the week WW and the day-of-week D .
minute	This function returns an integer representing the minute when the fraction of the expression is interpreted as a time according to the standard number interpretation.

3 Managing data in the app with Data manager

Function	Description
month	This function returns a dual value: a month name as defined in the environment variable MonthNames and an integer between 1-12. The month is calculated from the date interpretation of the expression, according to the standard number interpretation.
monthend	This function returns a value corresponding to a timestamp of the last millisecond of the last day of the month containing date . The default output format will be the DateFormat set in the script.
monthname	This function returns a display value showing the month (formatted according to the MonthNames script variable) and year with an underlying numeric value corresponding to a timestamp of the first millisecond of the first day of the month.
monthsend	This function returns a value corresponding to a timestamp of the last millisecond of the month, bi-month, quarter, tertial, or half-year containing a base date. It is also possible to find the timestamp for a previous or following time period.
monthsname	This function returns a display value representing the range of the months of the period (formatted according to the MonthNames script variable) as well as the year. The underlying numeric value corresponds to a timestamp of the first millisecond of the month, bi-month, quarter, tertial, or half-year containing a base date.
monthsstart	This function returns a value corresponding to the timestamp of the first millisecond of the month, bi-month, quarter, tertial, or half-year containing a base date. It is also possible to find the timestamp for a previous or following time period.
monthstart	This function returns a value corresponding to a timestamp of the first millisecond of the first day of the month containing date . The default output format will be the DateFormat set in the script.
networkdays	The networkdays function returns the number of working days (Monday-Friday) between and including start_date and end_date taking into account any optionally listed holiday .
now	This function returns a timestamp of the current time from the system clock. The default value is 1.
quarterend	This function returns a value corresponding to a timestamp of the last millisecond of the quarter containing date . The default output format will be the DateFormat set in the script.
quartername	This function returns a display value showing the months of the quarter (formatted according to the MonthNames script variable) and year with an underlying numeric value corresponding to a timestamp of the first millisecond of the first day of the quarter.

3 Managing data in the app with Data manager

Function	Description
quarterstart	This function returns a value corresponding to a timestamp of the first millisecond of the quarter containing date . The default output format will be the DateFormat set in the script.
second	This function returns an integer representing the second when the fraction of the expression is interpreted as a time according to the standard number interpretation.
timezone	This function returns the name of the current time zone, as defined in Windows.
today	This function returns the current date from the system clock.
UTC	Returns the current Coordinated Universal Time.
week	This function returns an integer representing the week number according to ISO 8601. The week number is calculated from the date interpretation of the expression, according to the standard number interpretation.
weekday	This function returns a dual value with: A day name as defined in the environment variable DayNames . An integer between 0-6 corresponding to the nominal day of the week (0-6).
weekend	This function returns a value corresponding to a timestamp of the last millisecond of the last day (Sunday) of the calendar week containing date . The default output format will be the DateFormat set in the script.
weekname	This function returns a value showing the year and week number with an underlying numeric value corresponding to a timestamp of the first millisecond of the first day of the week containing date .
weekstart	This function returns a value corresponding to a timestamp of the first millisecond of the first day (Monday) of the calendar week containing date . The default output format is the DateFormat set in the script.
weekyear	This function returns the year to which the week number belongs according to ISO 8601. The week number ranges between 1 and approximately 52.
year	This function returns an integer representing the year when the expression is interpreted as a date according to the standard number interpretation.
yearend	This function returns a value corresponding to a timestamp of the last millisecond of the last day of the year containing date . The default output format will be the DateFormat set in the script.
yearname	This function returns a four-digit year as display value with an underlying numeric value corresponding to a timestamp of the first millisecond of the first day of the year containing date .

3 Managing data in the app with Data manager

Function	Description
yearstart	This function returns a timestamp corresponding to the start of the first day of the year containing date . The default output format will be the DateFormat set in the script.
yeartodate	This function finds if the input timestamp falls within the year of the date the script was last loaded, and returns True if it does, False if it does not.

Formatting and interpretation functions that can be used in a calculated field expression

The formatting functions use the numeric value of the input expression, and convert this to a text value. In contrast, the interpretation functions do the opposite: they take string expressions and evaluate them as numbers, specifying the format of the resulting number. In both cases the output value is dual, with a text value and a numeric value.

For example, consider the differences in output between the **Date** and the **Date#** functions.

Date and the Date# functions

Function	Output (text)	Output (numeric)
Date#('20140831', 'YYYYMMDD')	20140831	41882
Date(41882, 'YYYY.MM.DD')	2014.08.31	41882

These functions are useful when your data contains date fields that are not interpreted as dates as the format does not correspond to the date format setting in Qlik Sense. In this case, it can be useful to nest the functions:

```
Date(Date#(DateInput, 'YYYYMMDD'), 'YYYY.MM.DD')
```

This will interpret the DateInput field according to the input format, YYYYMMDD, and return it in the format you want to use, YYYY.MM.DD.

Formatting and interpretation functions

Function	Description
Date	Date() formats an expression as a date using the format set in the system variables in the data load script, or the operating system, or a format string, if supplied.
Date#	Date# evaluates an expression as a date in the format specified in the second argument, if supplied.
Dual	Dual() combines a number and a string into a single record, such that the number representation of the record can be used for sorting and calculation purposes, while the string value can be used for display purposes.
Interval	Interval() formats a number as a time interval using the format in the system variables in the data load script, or the operating system, or a format string, if supplied.
Interval#	Interval#() evaluates a text expression as a time interval in the format set in the operating system, by default, or in the format specified in the second argument, if supplied.

3 Managing data in the app with Data manager

Function	Description
Money	Money() formats an expression numerically as a money value, in the format set in the system variables set in the data load script, or in the operating system, unless a format string is supplied, and optional decimal and thousands separators.
Money#	Money#() converts a text string to a money value, in the format set in the load script or the operating system, unless a format string is supplied. Custom decimal and thousand separator symbols are optional parameters.
Num	Num() formats an expression numerically in the number format set in the system variables in the data load script, or in the operating system, unless a format string is supplied, and optional decimal and thousands separators.
Num#	Num#() converts a text string to a numerical value, in the number format set in the data load script or the operating system. Custom decimal and thousand separator symbols are optional parameters.
Text	Text() forces the expression to be treated as text, even if a numeric interpretation is possible.
Time	Time() formats an expression as a time value, in the time format set in the system variables in the data load script, or in the operating system, unless a format string is supplied.
Time#	Time#() evaluates an expression as a time value, in the time format set in the data load script or the operating system, unless a format string is supplied.
Timestamp	TimeStamp() formats an expression as a date and time value, in the timestamp format set in the system variables in the data load script, or in the operating system, unless a format string is supplied.
Timestamp#	Timestamp#() evaluates an expression as a date and time value, in the timestamp format set in the data load script or the operating system, unless a format string is supplied.

Numeric functions that can be used in a calculated field expression

You can use these functions to round numeric values.

Numeric functions

Function	Description
ceil	Ceil() rounds up a number to the nearest multiple of the step shifted by the offset number.
div	Div() returns the integer part of the arithmetic division of the first argument by the second argument. Both parameters are interpreted as real numbers, that is, they do not have to be integers.
evens	Even() returns True (-1), if integer_number is an even integer or zero. It returns False (0), if integer_number is an odd integer, and NULL if integer_number is not an integer.
fabs	Fabs() returns the absolute value of x . The result is a positive number.
fact	Fact() returns the factorial of a positive integer x .
floor	Floor() rounds down a number to the nearest multiple of the step shifted by the offset number.

3 Managing data in the app with Data manager

Function	Description
fmod	fmod() is a generalized modulo function that returns the remainder part of the integer division of the first argument (the dividend) by the second argument (the divisor). The result is a real number. Both arguments are interpreted as real numbers, that is, they do not have to be integers.
frac	Frac() returns the fraction part of x .
mod	Mod() is a mathematical modulo function that returns the non-negative remainder of an integer division. The first argument is the dividend, the second argument is the divisor, Both arguments must be integer values.
odd	Odd() returns True (-1), if integer_number is an odd integer or zero. It returns False (0), if integer_number is an even integer, and NULL if integer_number is not an integer.
round	Round() returns the result of rounding a number up or down to the nearest multiple of step shifted by the offset number.
sign	Sign() returns 1, 0 or -1 depending on whether x is a positive number, 0, or a negative number.

Conditional functions that can be used in a calculated field expression

You can use these functions to evaluate a condition and then return different answers depending on the condition value.

Conditional functions

Function	Description
alt	The alt function returns the first of the parameters that has a valid number representation. If no such match is found, the last parameter will be returned. Any number of parameters can be used.
class	The class function assigns the first parameter to a class interval. The result is a dual value with $a \leq x < b$ as the textual value, where <i>a</i> and <i>b</i> are the upper and lower limits of the bin, and the lower bound as numeric value.
if	The if function returns a value depending on whether the condition provided with the function evaluates as True or False.
match	The match function compares the first parameter with all the following ones and returns the numeric location of the expressions that match. The comparison is case sensitive.
mixmatch	The mixmatch function compares the first parameter with all the following ones and returns the numeric location of the expressions that match. The comparison is case insensitive.
pick	The pick function returns the <i>n</i> :th expression in the list.
wildmatch	The wildmatch function compares the first parameter with all the following ones and returns the number of the expression that matches. It permits the use of wildcard characters (* and ?) in the comparison strings. * matches any sequence of characters. ? matches any single character.

3 Managing data in the app with Data manager

NULL functions that can be used in a calculated field expression

You can use these functions to return or detect null values.

NULL functions

Function	Description
Null	The Null function returns a NULL value.
IsNull	The IsNull function tests if the value of an expression is NULL and if so, returns -1 (True), otherwise 0 (False).

Mathematical functions that can be used in a calculated field expression

You can use these functions for mathematical calculations.

Mathematical functions

Function	Description
----------	-------------

e	The function returns the base of the natural logarithms, e (2.71828...).
----------	--

rand	The function returns a random number between 0 and 1. This can be used to create sample data.
-------------	---

Exponential and Logarithmic functions that can be used in a calculated field expression

You can use these functions for exponential and logarithmic calculations.

Exponential and Logarithmic functions

Function	Description
exp	The natural exponential function, e^x , using the natural logarithm e as base. The result is a positive number.
log	The natural logarithm of x . The function is only defined if $x > 0$. The result is a number.
log10	The common logarithm (base 10) of x . The function is only defined if $x > 0$. The result is a number.
pow	Returns x to the power of y . The result is a number.
sqr	x squared (x to the power of 2). The result is a number.
sqrt	Square root of x . The function is only defined if $x \geq 0$. The result is a positive number.

Distribution functions that can be used in a calculated field expression

You can use these functions for statistical distribution calculations.

Distribution functions

Function	Description
CHIDIST	CHIDIST() returns the one-tailed probability of the χ^2 distribution. The χ^2 distribution is associated with a χ^2 test.

3 Managing data in the app with Data manager

Function	Description
CHIINV	CHIINV() returns the inverse of the one-tailed probability of the chi ² distribution.
FDIST	FDIST() returns the F-probability distribution.
FINV	FINV() returns the inverse of the F-probability distribution.
NORMDIST	NORMDIST() returns the cumulative normal distribution for the specified mean and standard deviation. If mean = 0 and standard_dev = 1, the function returns the standard normal distribution.
NORMINV	NORMINV() returns the inverse of the normal cumulative distribution for the specified mean and standard deviation.
TDIST	TDIST() returns the probability for the Student's t-distribution where a numeric value is a calculated value of t for which the probability is to be computed.
TINV	TINV() returns the t-value of the Student's t-distribution as a function of the probability and the degrees of freedom.

Geospatial functions that can used in a calculated field expression

You can use this function to handle geospatial data.

Geospatial functions

Function	Description
GeoMakePoint	GeoMakePoint() is used in scripts and chart expressions to create and tag a point with latitude and longitude.

Color functions that can used in a calculated field expression

You can use these functions for setting and evaluating color properties.

Color functions

Function	Description
ARGB	ARGB() is used in expressions to set or evaluate the color properties of a chart object, where the color is defined by a red component r , a green component g , and a blue component b , with an alpha factor (opacity) of alpha .
HSL	HSL() is used in expressions to set or evaluate the color properties of a chart object, where the color is defined by values of hue , saturation , and luminosity between 0 and 1.
RGB	RGB() is used in expressions to set or evaluate the color properties of a chart object, where the color is defined by a red component r , a green component g , and a blue component b with values between 0 and 255.

Logical functions that can used in a calculated field expression

You can use these functions for handling logical operations.

3 Managing data in the app with Data manager

Logical functions

Function	Description
IsNum	Returns -1 (True) if the expression can be interpreted as a number, otherwise 0 (False).
IsText	Returns -1 (True) if the expression has a text representation, otherwise 0 (False).

System functions that can be used in a calculated field expression


You can use these functions for accessing system ,device, and Qlik Sense app properties.

System functions

Function	Description
OSUser	This function returns a string containing the name of the user that is currently connected. It can be used in both the data load script and in a chart expression.
ReloadTime	This function returns a timestamp for when the last data load finished. It can be used in both the data load script and in a chart expression.

Changing field types

When data is added, Qlik Sense interprets the field type of each field. The following field types are currently supported:

-  **General**
-  **Date**
-  **Timestamp**
-  **Geo data**


If the data was not interpreted correctly, you can adjust the field type. You can also change the input and display format of a data or timestamp field.

To open the table editor, click  on the data table you want to edit.


It is not possible to change field type or display format of fields in some cases.

- The table is the result of concatenating two or more tables.
- The field is already recognized as a date or a timestamp.

Making sure a date or timestamp field is recognized correctly

If a date or timestamp field is not recognized as a date or a timestamp, that is, it is marked with  **General**, you can make sure it is interpreted correctly.

Do the following:

1. Click on  above the field heading.
The data format dialog opens.

3 Managing data in the app with Data manager

2. Change **Field type** to **Date** or **Timestamp**.
3. Change the format string in **Input format** to interpret the date correctly. You can use a prepared format from the drop down list, or write your own format string.





It is not possible to use a single quote in the format string.

4. If you want to use a display format other than the default format in your app, write or select a format string in **Display format**.
If you leave it empty, the app default display format is used.

Changing the display format of a date or timestamp field

Each app has default display formats for date and timestamp fields. You can change the display format for an individual date or timestamp field.


Do the following:

1. Click on  (date) or  (timestamp) above the field heading.
The data format dialog opens.
2. Change the format string in **Display format**. Either use a prepared format from the drop down list, or write your own format string..

Changing a field type to geographical data

If a field with values such as city and country names or ISO symbols is not recognized as geographical data, you can change the field type to **Geo data**.

Do the following:

1. Click on  above the field heading.
The data format dialog opens.
2. Select **Geo data** from the **Field type** drop-down menu.
3. Select the type of geographical data from the **Geo data** drop-down menu.
The options are **City**, **Country**, **Country code ISO2**, and **Country code ISO3**. ISO2 and ISO3 are from ISO 3166, the International Standards Organization's codes for countries. ISO2 contains two-character codes, and ISO3 contains three-character codes. For example, the codes for Sweden are SE and SWE. When assigning an ISO code, be sure that the values in the field match the code you assign. If you assign ISO3 to a field with two-character code values, location coordinates will not be assigned correctly.
4. For **City** data, select the related field in the table that contains geographical data for countries. There may be only one related country field, but it is possible to have multiple fields with geographical data for countries. For example, one field may have long names such as France and another field that designates countries by country codes such as ISO2. But the fields appear in the **Associated country** list only if they have been designated as **Geo data**.
If no field has been designated as a **Geo data** country field, then the **Associated country** list does not appear when you designate a field as **City**.

3 Managing data in the app with Data manager

When a field is assigned the **Geo data** field type, either by the user or automatically by Qlik Sense, a field containing geographical coordinates, either point or polygon data, is associated with it. The associated fields containing the coordinates are visible in the **Data model viewer**. These coordinates are required for apps that use **Map** objects.

Fields that contain geographical information in the form of names or codes, such as postal areas, cannot be used for mapping unless they are designated as **Geo data** fields.

The fields assigned the **Geo data** type continue to hold string values, such as Mexico and MX, but when they are used in a **Map** object, the mapping coordinates come from the fields containing the point or polygon data.

Hiding fields from analysis

You can hide data fields that you do not want available when building visualizations in sheet view or in insight advisor. You might, for example, have fields that are only used for calculating another field. You can hide these fields in **Data manager** so that they are not available in the assets panel of sheets or from insight advisor but remain available in **Data manager** and **Data load editor**.


When hiding a field, all existing relationships the field has, such as associations or use in calculations, will be maintained. If a field is currently in use, such as in a master item or in an existing chart, it will continue to be available there, but will not be available for use in new master items or visualizations.




*You can view all your hidden fields in **Data manager** by going into **Data load editor**, and opening the auto-generated section. All hidden fields will be listed as TAG FIELD <field name> WITH '\$hidden';.*

Hiding a field from analysis

Do the following:

1. Click on  above the field heading.
2. Click **Hide in Analysis**.

The field is now hidden in sheet view and in insight advisor. Hidden fields have  added above the field heading.

Showing a hidden field

Do the following:

1. Click on  or  above the field heading.
2. Click **Show in Analysis**.

The field is now available in sheet view and in insight advisor.

Assessing table field data before loading data

To examine your data for potential quality issues such as null or outlier before you load it into Qlik Sense, you view a summary of the data using the **Summary** data profiling card. In addition, the **Summary** card enables you to view different possible data interpretations, such as viewing the field's data as a dimension or measure.


You access the **Summary** card by editing a table in **Data manager** and selecting a table field. When a field is selected in the table editor, Qlik Sense examines the data type, metadata, and values present. The field is then categorized as either a dimension, measure, or temporal field, and an analysis is presented in the **Summary** card. Fields whose data can be categorized as either a dimension or measure can switch preview to display them as a dimension or measure. How a data field is categorized in the **Summary** card does not affect how you can use it in Qlik Sense visualizations, but it does determine what transformation options are available for the data field in other data profiling cards.

A summary of data using the Summary data profiling card.


Gross Margin Summary (Measure) ?

Distinct values: **225** Total values: **713**

Preview of distinct value distribution



Value Range



The **Summary** card provides the following information:

- **Distinct values:** The number of distinct values in the field.
- **Total values:** The number of values in the field.
- **Preview of distinct value distribution:** In fields with more than 20 distinct values, only the 20 distinct values with the highest count display. All other values are grouped into a single value on the chart. If all values are distinct, no bar chart will display.

3 Managing data in the app with Data manager

- **Value Range:** (Measure and Temporal only) For a measure field, the **Value Range** is a chart showing the Min, Median, Average, and Max values for the field. For a temporal field, the **Value Range** is the time period covered by the field's data.
- **Null values:** The number of null values in the data. This visualization only displays if there are null values in the field.
- **Mixed values:** The number of text value in a field that contains both text and numeric values. This visualization only displays if there are mixed values in the field.


Depending on how a field is categorized in the **Summary** card, it can be modified in other data profiling cards. Fields set as measures can have grouped values created from the field using the **Bucket** card. For more information, see *Grouping measure data into ranges (page 75)*.

Fields set as dimensions can have:

- Distinct values replaced with other value using the **Replace** card.
For more information, see *Replacing field values in a table (page 69)*.
- Distinct values set as null values using the **Set nulls** card.
For more information, see *Setting field values as null in a table (page 71)*.
- A custom order applied to the values using the **Order** card.
For more information, see *Customizing the order of dimension values (page 72)*.
- Field data split into new table fields using the **Split** card.
For more information, see *Splitting a field in a table (page 73)*.

Accessing the **Summary** card

Do the following:

1. In **Data manager**, select a table and click .
2. Select a table field.

The Summary card appears.

Changing the data category of a field

Qlik Sense permits changing the data category of a field, provided the data in the field can be categorized in multiple categories.

Do the following:

- In the **Summary** card, click ▼ and select a different data category.

Replacing field values in a table

The **Replace** card enables you to replace different field values in your tables with other values. In a data set, different terms might be used for the same object or concept. For example, the full name of a country and its abbreviation might both be found in the same data set. The **Replace** card enables you to ensure that these

3 Managing data in the app with Data manager

instances are treated as a single distinct value rather than different distinct values. For example, in a field that contains country data, you could replace *U.S*, *US*, and *U.S.A* with *USA*. You can also use the **Replace** card to change individual values, such as when a name in a data set needs to be changed.


You can set replace values in fields that contain up to a maximum of 5,000 distinct values.

In addition, calculated fields ignore replacement values and will use the original values instead.

The **Replace** card consists of two sections: **Distinct values** and **Replacement value**. **Distinct values** lists all distinct values and any replacement values. **Replacement value** contains a field to enter the replacement term and a list of values selected to be replaced. You replace field values by selecting distinct values, entering the replacement value, and applying the replacement. You can also select a replacement value in **Distinct values** and edit it to change the values being replaced or the value being used for the replacement. You can add or edit multiple replacement values before applying the replacements.

Replacing values

Do the following:

1. In **Data manager**, select a table and click .
2. Select a dimension field.
3. In the data profiling card, click the **Replace** tab.
4. Under **Distinct values**, select the distinct values you want to replace.
5. Under **Replacement value**, enter the new value for these distinct values.
6. Click **Replace**.

Editing a replacement value

You can add more distinct values to a replacement value, remove distinct values from a replacement value, or change the replacement value.


Adding distinct values to a replacement value

Do the following:

1. In the **Replace** card, under **Distinct values**, select a replacement value.
2. Under **Distinct values**, select the distinct values you want to add.
3. Click **Replace**.

Removing distinct values from a replacement value

Do the following:

1. In the **Replace** card, under **Distinct values**, select a replacement value.
2. After each distinct value you want to remove, click .
3. Click **Replace**.

Changing a replacement value

Do the following:

1. In the **Replace** card, under **Distinct values**, select a replacement value.
2. Under **Replacement**, enter a new value.
3. Click **Replace**.

Deleting a replacement value

Replacement values can be deleted by removing all associated distinct values.

Do the following:

1. In the **Replace** card, under **Distinct values**, select a replacement value.
2. Click **Remove All**.
3. Click **Replace**.

Setting field values as null in a table

In the **Set nulls** card, you select distinct values from a dimension field to be treated as null values in Qlik Sense. For example, if your source data includes representations such as X for nulls, you can use the **Set nulls** card to set that value to be treated as a null value in Qlik Sense. If your table contains fields with empty spaces, you can set these as null values using the **Set nulls** card. You can also use the **Set nulls** card to clean your data of unwanted values by setting these values as null.


If you want to use a specific value as your null value, you can replace the default null value, - (**Null**), using the **Replace** card. For more information, see *Replacing field values in a table (page 69)*.

You can set field values as null in fields that contain up to a maximum of 5,000 distinct values.

The **Set nulls** card consists of two sections, **Distinct values** and **Manual null values**. When you select values from **Distinct values**, they are added to **Manual null values**. When you apply the null values, all instances of the selected values are set to null in the field's data. You can restore individual or all values set as null.


Setting field values as null


Do the following:

1. In **Data manager**, select a table and click .
2. Select a field.
3. In the data profiling card, click the **Set nulls** card.
4. Under **Distinct values**, select the values you want set as null.
5. Click **Set null values**.

Restoring values manually set as null

Do the following:

1. In **Data manager**, select a table and click .
2. Select a field.
3. In the data profiling card, click the **Set nulls** card.

4. In the **Set nulls** card, under **Manual null values**, do one of the following:
 - Click  after the values you no longer want set as null.
 - Click **Remove All** to restore all values set as null.
5. Click **Set null values**.

Customizing the order of dimension values

Custom ordering enables you to set the order of dimension values in visualizations. While alphabetical or numerical order may work for many visualizations, in some instances, an alternative order is useful. For example, in a field containing cities, you could set a custom order so that in your charts cities are listed in order of east to west. You can define custom ordering for fields that meet the following requirements:


- Fields must be set as dimensions in the **Summary** card.
- Fields must contain up to a maximum of 5,000 distinct values.
It is recommended that your field have up to a maximum of 25 values, but you can choose to use the **Order** card with fields containing up to 5,000 distinct values.

The **Order** card consist of two sections: **Current Order** and **Preview of order**. **Current Order** displays all the distinct values from the dimension. By default, the distinct values are organized by load order. You set your custom order by dragging the values in **Current Order** into the desired order. **Preview of order** is a bar chart that displays a count of values in each distinct value, organized by the current order.

A custom order overrides all other sorting options available in Qlik Sense visualizations except for sorting by load order. If you require alphabetical or numeric ordering for this field, you must remove the custom order by resetting the order.

Changing the order of values in a dimension field

Do the following:

1. In **Data manager**, select a table and click .
2. Select a dimension field.
3. In the data profiling card, click the **Order** tab.
4. Click and drag the values into the new order.
5. If you want to cancel your current ordering, click **Cancel**.




***Cancel** is only available to new custom orders. To cancel your changes if you are changing the order of values in an existing custom order, select a different field in the table and then select this field again.*

6. Click **Reorder**.

Resetting the order of values in a dimension field

Do the following:

1. In **Data manager**, select a table and click .
2. Select a reordered dimension field.
3. In the data profiling card, click the **Order** tab.
4. Click **Reset**.
5. Click **OK**.

Splitting a field in a table

The **Split** card enables you to create new fields using data from an existing field. You could, for example, split a field that contains an address to create a new field that contains only a zip or postal code. This lets you quickly create new fields containing segments of existing data. The **Split** card can create new fields from table fields that meet the following requirements:

- Fields must be set as dimensions in the **Summary** card.



*Table fields containing date and time information are automatically split into date fields when their tables are prepared in **Data manager** and do not require the **Split** card.*

The **Split** card consists of an input field containing a template value and a preview of the new fields with their values. By default, the template value is the first value in numerical order from a field, but you can select other values from the source field to use as the template. You should select a value that is representative of all values in the table. Splitting using an outlier value as a template may impact the quality of the new fields.

Fields are split by inserting split markers into the template value where you want to split the field. Split markers are added by selecting a point in the sample field where you want to add a split marker, adjusting your selection, and then selecting to split by instance or position. The **Split** card may automatically add recommended split markers to your template value.

Instances are occurrences of a selected delimiter, such as the character @ or a space between words. Positions of instance split markers are relative to either:

- The start of the value, such as the first instance of @ in a value.
- Their position to the right of another instance or position split marker, such as the first instance of . after @.

If you remove the instance to which another instance is relative, the other instance's position adjusts to the same position relative to the next instance of a different delimiter set as a split marker or the start of the value. You can split a field on up to 9 delimiters.



*The **Split** card splits values using the characters specified as the split markers. If the data has variances in how these characters are composed, such as accented characters, those variances will not be included in the split.*

Positions are locations in the field value, such as after the first four characters. Positions are relative to either:

3 Managing data in the app with Data manager

- The start of the value.
- Their position to the right of an instance split marker, such as the second character after an instance split marker.

If you removed an instance that has a position to its right, the position moves to the same position relative to the start of the value or the first next instance split marker to the left or the start of the value.

The field preview updates as you add split markers, showing the new fields and their data. You can rename the new fields in the field preview. You can also select to include or exclude split fields from your table before you apply the split. When you apply a split, the fields you selected in the field preview are added to your table.


A field can be split multiple times.

To split a field into new fields, do the following:

1. Access the **Split** card.
2. Set split markers in the input field.
3. You can also perform these optional tasks:
 - Remove split markers.
 - Select which new fields you want to be added to the table.
 - Rename the new fields.
4. Create the new fields.

Accessing the Splitting card

Do the following:

1. In **Data manager**, select a table and click .
2. Select a dimension field.
3. In the data profiling card, click the **Split** tab.

Inserting split markers

You insert split markers by clicking positions in the sample value and selecting the kind of split you want to apply.



You can change the sample value displayed in the input field by clicking ▼ and selecting a different value.

Do the following:


1. In the **Split** card, click the position in the sample value where you want to add split markers.
Clicking selects all of the template value up to any other split markers.
Double-clicking selects the insert point of your cursor in the template value.
2. Adjust your selection by clicking and dragging the selection tabs or highlighting the section you want to select.
3. Click the button that corresponds to the kind of split you want applied:

- **This instance:** The field splits by the selected instance of the delimiter.
- **All instances:** The field splits by all instances of the delimiter.
- **These positions:** The field splits on either side of the selection.
- **This position:** The field splits at this position.

The split marker is inserted into the template value.

Removing split markers

Do the following:

- In the **Split** card, do one of the following:
 - To remove a single split marker, click  above the split marker you want to remove.
 - To remove all split marker, click **Reset**.
Recommended split markers will not be removed and any removed recommended split markers will be added back to the template. You must remove these individually.

Selecting the fields to add to your table

You can select which fields created by the **Split** card to include or exclude from your table. By default, all split fields are included.

Do the following:

- In the **Split** card, in the field preview, do one of the following:
 - To include a field, select the field column checkbox.
 - To exclude a field, clear the field column checkbox.

Renaming new fields

Do the following:

- In the **Split** card, in a field header in the field preview, enter a new field name.

Creating new fields

Do the following:

- In the **Split** card, click **Split**.

Grouping measure data into ranges

Measures can provide additional insight for analysis when grouped into ranges and used as dimensions in visualizations. For example, if a field had customer ages, you could group values into age ranges. The **Bucket** data profiling card enables the grouping of field data into ranges, creating a new field with the specified groupings. You access the **Bucket** card by editing a table in **Data manager** and selecting a table field that is set as a measure in the **Summary** card. You can group measures into ranges with fields that meet the following requirements:

3 Managing data in the app with Data manager

- Fields must be classified as a measure in the **Summary** card
- Fields must contain at least 10 distinct values.
The **Bucket** card is not recommended with fields containing a low range of values, but you can choose to use the **Bucket** card with non-recommended fields.
- Values in the field cannot contain more than 20% mixed types.
The **Bucket** card is not recommended with fields containing mixed types of values, but you can choose to use the **Bucket** card with non-recommended fields.



*Bucket fields created by the **Bucket** card are categorized as dimensions in the **Summary** card. Bucket fields cannot have a custom order applied to them using the **Order** card, however. Bucket fields cannot be used in calculated fields.*

The **Bucket** card provides a suggested number of groupings, a preview of the groupings, and a slider bar containing your groupings that enables modifying each bucket's name and value range. You can modify the suggested number of groupings by entering a new number into the **Bucket** field. Qlik Sense supports a maximum of 20 buckets and a minimum of 2 buckets. New buckets are added to the right of your bucket range while buckets are removed from right to left. If you have not modified any individual buckets, each bucket's range is adjusted to fit evenly within the value range of the field. If you change the number of buckets after having modified an individual bucket, new buckets are added to the right of your bucket range and are given a range of values equal to the size of the second rightmost bucket.

The **Preview of data buckets** bar chart gives overview of the data in your buckets, with a count of the number of distinct values in each bucket. The chart is updated as you alter your buckets. If a bucket has no values in it, it will have no bar in the chart.

The **Bucket** slider bar enables you to edit your buckets. By clicking a bucket segment, you can set that bucket's range, change the bucket's name, or remove the bucket entirely. Hovering your cursor over the bucket brings up the bucket's name and range of values. By default, buckets are named with the bucket's value range expressed as an interval notation. Buckets ranges include values between the starting value and up to but excluding the ending value.

When you adjust a single bucket's value range, Qlik Sense shifts the values of all buckets, ensuring there are no gaps and overlaps while respecting the existing quantitative ranges of the other buckets as much as possible. The leftmost bucket always has no lower boundary and the rightmost bucket has no upper boundary. This allows them to capture any values that might fall outside the set ranges of all your buckets. Modifying the lower range of a bucket will alter the ranges of the buckets to the right, modifying the upper range of a bucket will modify the buckets to the left.

When you create buckets from a field, a new field is generated containing all the buckets assigned to the rows with their corresponding measure values from the source field. By default, it will be named <field> (Bucketed). This field can be renamed, associated, sorted, and deleted like other table fields. You can edit the bucketing in the generated bucket field by selecting the field in the table and modifying the bucketing options. You can create multiple bucket fields from the same source measure field.


To group measure values into ranges, do the following:

3 Managing data in the app with Data manager

1. Access the **Bucket** card.
2. Optionally, change the number of buckets.
3. Optionally, modify individual buckets.
4. Optionally, reset the manual adjustments to buckets to the default settings.
5. Create the bucketed field.

Accessing the Buckets card

Do the following:

1. In **Data manager**, select a table and click  .
2. If you are creating new groups of value ranges, select a field.
3. If you are editing an existing buckets field, select a buckets field

Changing the number of buckets

Changing the number of buckets adds buckets to or removes buckets from the end of the range.

Do the following:

- In the **Bucket** card, enter a new number before **Buckets**.

Modifying a bucket

You can rename a bucket, adjust a bucket's value range, or delete a bucket.



If you are modifying buckets, it is recommended to modify buckets in order from the leftmost bucket segment to the rightmost bucket segment.

Renaming a bucket

Do the following:

1. In the **Bucket** card, click the bucket segment.
2. In the name field, enter a new name.
3. To apply your changes, click anywhere outside the bucket segment.

Adjusting a bucket's range of values



*If you are trying to use decimal values, you must enter them manually in the **From** and **Up To** fields.*


Do the following:

3 Managing data in the app with Data manager

1. In the **Bucket** card, click the bucket segment.
2. Do one of the following:
 - After **From** and **Up To**, enter new values to set the range of the bucket.
A bucket includes the values between the value in **From** and up to but excluding the value in **Up To**.
 - Adjust the segment's sliders to set the range of the bucket.
3. To apply your changes ,click anywhere outside the bucket segment.

Deleting a bucket

Do the following:

1. In the **Bucket** card, click the bucket's segment.
2. Click  .

Resetting a field's buckets

Resetting buckets enables you to restore the bucket card settings to their default state. If this is a measure field, the default state is the Qlik Sense recommended bucketing. If this is a bucket field, the default state is the field's state after the last time **Create buckets** was clicked.

Do the following:

- In the **Bucket** card, click **Reset to default**.

Creating a bucket field

Do the following:


1. In the **Bucket** card, click **Create buckets**.
2. Click **OK**.

The new field containing your data grouping is added to your table.

Unpivoting crosstab data in the data manager

A crosstab is a common type of table featuring a matrix of values between two orthogonal lists of header data. It is usually not the optimal data format if you want to associate the data to other data tables. This topic describes how you can unpivot data loaded in crosstab format, that is, transpose parts of it into rows using the data manager.

Unpivot data loaded in crosstab format transpose parts of it into rows.



Year	Europe	RoW
2007	234	567
2008	345	534

Year	Region	Sales
2007	Europe	234
2007	RoW	567
2008	Europe	345
2008	RoW	534

What's a crosstab?

A crosstab contains a number of qualifying columns, which should be read in a straightforward way, and a matrix of values. In this case there is one qualifying column, Year, and a matrix of sales data per month.

Crosstab

Year	Jan	Feb	Mar	Apr	May	Jun
2008	45	65	78	12	78	22
2009	11	23	22	22	45	85
2010	65	56	22	79	12	56
2011	45	24	32	78	55	15
2012	45	56	35	78	68	82

If this table is simply loaded into Qlik Sense, the result will be one field for *Year* and one field for each of the months. This is generally not what you would like to have. You would probably prefer to have three fields generated:

- The qualifying field, in this case *Year*, marked with green in the table above.
- The attribute field, in this case represented by the month names Jan - Jun marked with yellow. This field can suitably be named *Month*.
- The data field, marked with blue. In this case they represent sales data, so this can suitably be named *Sales*.

This can be achieved by using the Unpivot option in the data manager table editor, and selecting the fields Jan - Jun. This creates the following table:


Unpivoted table

Year	Month	Sales
2008	Jan	45
2008	Feb	65
2008	Mar	78
2008	Apr	12
2008	May	78
2008	Jun	22
2009	Jan	11
2009	Feb	23
...

Unpivoting a crosstab table into a flat table

Do the following:

3 Managing data in the app with Data manager

1. Add a data file in crosstab format to your app.
2. Click  on the table in the data manager to open the table editor.
3. Click **Unpivot**.
4. Select the fields you want to transpose into rows. You need to have at least one qualifying field that is not unpivoted. There are two ways to make the selections.
 - Click on the field headers of the fields you want to transpose. Do not select the fields you want to keep as qualifying fields.
 - Click on the field headers of the fields you want to keep as qualifying fields, and then select **Invert selections** from the field menu. This is the easiest way to do it if you have a large number of fields to transpose.
5. Click **Apply unpivoting**
The selected data is now transposed to rows with two fields, Tablename.**Attribute field** and Tablename.**Data field**.
6. Rename **Attribute field** to something meaningful, in the example above, *Month*.
7. Rename **Data field** to something meaningful, in the example above, *Sales*.

You have now unpivoted the crosstable to a flat format, which will make it easier when you want to associate it to other data in the app.


Reverting to the original crosstab table

You can revert to the original crosstab format of your data source by clicking **Undo unpivot** in the table editor. If you created any associations to other data in the app, the associations will be deleted.

3.13 Concatenating tables in **Data manager**

Concatenation combines two tables into a single table with combined fields. It consolidates content, reducing the number of separate tables and fields that share content. Tables in **Data manager** can be automatically or forcibly concatenated.

Concatenated table in Tables view and Associations view.



Automatically concatenating tables

Tables are automatically concatenated in **Data manager** when Qlik Sense detects that one or more added tables have both the same number of fields and identical field names as another table. When this happens, Qlik Sense automatically concatenates the tables into a single table. Automatically concatenated tables can be split if

3 Managing data in the app with Data manager

they were concatenated erroneously or if you do not want them concatenated. Automatically concatenated tables can be forcibly concatenated to other tables.

Automatically concatenated tables have the following restrictions:

- You cannot change field categories.
- You cannot unpivot an automatically concatenated table.
- You cannot add or remove data with **Select data from source**.

Forcing concatenation between tables

Concatenation can be forced between tables in **Data manager** using the **Concatenate tables** pane, even if they do not entirely share the same fields or data. Forced concatenation enables you to control mapping and exclude fields from the final concatenated table. Forced concatenation supports additional transformations. Using forced concatenation, you can:

- Concatenate a concatenated table with another table.
- Concatenate an unpivoted table with another table. Forcibly concatenated tables can be unpivoted.
- Concatenate tables with calculated fields. Calculated fields can be concatenated to other fields in a forced concatenation. Calculated fields can be added to forcibly concatenated tables.



Forcibly concatenated tables have the following restrictions:



- Forced concatenation requires at least one field from each table be included in the concatenated table, although they need not be mapped together.
- Date fields cannot be formatted after concatenation. Date fields must have the same format applied to them before concatenation. Concatenated date fields use the default time format set with **DateFormat** in the **Data load editor**.
- You cannot change field categories after concatenation.
- Calculated fields that refer to a field mapped to another field in a concatenated table will only contain data for the original field rather than the combined data in the concatenated field. Calculated fields created after two tables are concatenated that refer to a field in the concatenated table will use all data in that field.
- You cannot add or remove data from a concatenated table with **Select data from source**. You can, however, remove fields by clicking **Add data**, selecting the source table, and then excluding the fields. Null values are added for the removed field's data.

The **Concatenate tables** pane is accessed by clicking **...** in **Data manager**, clicking **Concatenate tables**, and selecting two tables. When tables are selected in **Concatenate tables**, Qlik Sense analyzes the fields and automatically maps any fields together that match. If there are no clear matches, fields are left unmapped. When the concatenation is applied, mapped fields are combined in the concatenated table, while unmapped fields are included as individual fields with null values for the rows where there is no corresponding value.

The first table selected in **Concatenate tables** is set as the primary table, the table to which the other table is concatenated. The concatenated table uses the table and field names from the primary table unless these are

3 Managing data in the app with Data manager

manually renamed. You can change which table is the primary table with the  button. **Concatenate tables** arranges fields in two rows, with the primary table fields in the top row and the secondary table fields on the bottom row. You can swap the primary and secondary tables with the  button.

You can use **Edit mappings** to change the default mapping and select which fields to map, leave unmapped, or to exclude from the concatenated table. **Edit mappings** contains a drag and drop interface for editing mappings and the **Fields** pane, which lists all table fields. Fields can be mapped by dragging them beneath a primary table field. Fields can be added as a new unmapped field by  beside the field in the **Fields** pane or dragging them into the top row of fields. Unmapped fields are marked with  in the **Fields** pane. Fields removed from the concatenated table are not included in the table and are not available for use in Qlik Sense after concatenation is applied to the table.




Once mappings are applied and the tables are concatenated, you cannot edit them, but they can be removed from the tables by splitting the concatenated table, which restores the tables to their original state.

To forcibly concatenate tables in **Data manager**, do the following:

1. Select tables for concatenation.
2. Optionally, rename the concatenated table and the field names.
3. Optionally, edit the concatenation mappings.
4. Concatenate the tables.

Selecting tables for concatenation

Do the following:

1. In **Data manager**, click  in the bottom row.
2. Click **Concatenate tables**.
The **Concatenate tables** pane opens.
3. Select two tables.
The fields of both tables will be mapped or left unmapped in the **Concatenate tables** pane.
4. To preview a sample of unique values in each field, click .
5. To switch the primary and secondary tables, click .

Renaming the table and field names

Do the following:



1. In the **Concatenate tables** pane, in the table name field, enter a new table name.
2. In a field name field, enter a new field name.

Editing concatenation mappings

Do the following:

1. In the **Concatenate tables** pane, click **Edit mappings**.
2. To map two fields, click and drag a table field under a primary table field.

3 Managing data in the app with Data manager

3. To add a new unmapped field, click and drag a table field into the upper row of fields.
4. To remove a field from the concatenated table, in the field click .
5. To return a removed field back to the table, click in the **Fields** pane, click  beside the field.
6. Click **Edit mappings** to close **Edit mappings**.

Concatenating tables

Do the following:



- In the **Concatenate tables** pane, click **Apply**.

The tables are now concatenated.

Splitting concatenated tables


In cases where concatenation is no longer needed, such as when Qlik Sense has performed an unwanted automatic concatenation, you can split the concatenated tables into their source tables.



*Splitting a concatenated table will remove any associations the concatenated table had as well as any associations the primary and secondary tables had with each other. If you want to preserve your associations while splitting concatenated tables, click  to undo the concatenation instead of splitting the table. You cannot use  to undo concatenation after you load data in **Data manager**.*

Splitting an automatically concatenated table

Do the following:

1. Select the concatenated table.
2. Click .
3. Select the tables to split from the concatenated table.
4. Click **Split**.

The table is now split into its source tables and all fields in the source tables are qualified. Qualified fields are renamed with the table name followed by the field name, separated by a period punctuation mark (the character “.”).

Example:


Table1 and Table2 both contain the fields Field1 and Field2. When you add them in **Data manager**, they are concatenated to a table called Table1-Table2 with the same fields, Field1 and Field2.

If you split Table1-Table2, the result is two tables:

- Table1 with fields Table1.Field1 and Table1.Field2
- Table2 with fields Table2.Field1 and Table2.Field2

Splitting a forcibly concatenated table

Do the following:

1. Select the concatenated table.
2. Click  .

The table is now split into its source tables. All fields in the source tables and their fields have their pre-concatenation names. Splitting a concatenated table only splits one level of concatenation, so that any concatenated tables that were part of the split concatenated table have their own concatenation preserved.

3.14 Viewing table and field transformation details in **Data manager**

You can view the operations and transformations performed on tables and fields in **Data manager** using the **Details** dialog. The **Details** dialog is available in the **Associations** and **Table** views for tables and in the data table editor for fields.

Details displays the current operations and transformations made to the selected table or field, in the order they are applied in the generated data load script. This enables you to easily see the source of a table or field, the current changes that have been made, and the sequence in which the changes have been applied. You can use **Details**, for example, to easily see which tables were concatenated or if a field was reordered.

The information displayed in **Details** varies depending if you are viewing a table or field. Table **Details** displays:


- Source tables for the selected table.
- Transformations used on the table, such as unpivoting and concatenation.

Field **Details** displays:

- Source tables and fields for the selected field.
- Field type changes.
- Transformations used on the fields, such as from the data profiling card or from concatenation.

Viewing table details



Do the following:

- In **Data manager**, select a table, click  , and click **View details**.

The **Details** dialog opens.

Viewing field details

Do the following:

1. In **Data manager**, select a table and click .
2. Click  above a field heading and click **View details**.

The **Details** dialog opens.

3.15 Step-by-step - Combining tables using forced concatenation

This step-by-step walkthrough shows how you can use forced concatenation to combine two similar data tables.

Forced concatenation can be used to clean up your data before you use it for analysis in a sheet. You can concatenate two tables into one table. You can also add another table later, for example if you initially add a table from June, and then later want to add a second table from July.

Concatenation at a glance

- Tables are automatically concatenated in Data manager when Qlik Sense detects that one or more added tables have both the same number of fields and identical field names as another table. In this case, you can split the tables if needed.
- Two tables can be force concatenated when tables do not entirely share the same fields or data. Only two tables can be force concatenated. To concatenate three tables, for example, concatenate the first two tables into one table. Concatenate the third table to that created table.
- Tables that are not similar enough will not automatically be concatenated. You also will not be able to forcibly concatenate them. In this case, the fields in the table should instead be associated in the Data manager.

Walkthrough - Forced concatenation

These are the tasks required to complete the walkthrough:

1. Prepare the data tables
2. Add data tables to an app
3. Concatenate and load data tables into an app
4. A step further - adding a new table and concatenating the data fields

Prerequisites

You should know how to create an app in Qlik Sense.

Prepare the data tables

We have supplied some sample data for you to use to demonstrate forced concatenation. You can also use your own data, provided the fields and data are mostly the same in your two tables.

For example, here is the header and first row of the data that we supplied below. It has been pasted into two Excel tables. Note the differences in the fields.

3 Managing data in the app with Data manager

	A	B	C	D	E	F	G	H
1	SalesOrderID	SalesOrderDetailID	TrackingNumber	OrderQty	PID	SpecialOfferID	UnitPrice	Modified Date
2	43659	1	4911-403C-98		776	1	1822.4946	7/1/2013 0:00

	A	B	C	D	E	F	G	H
1	SalesOrderID	SalesOrderDetailID	TrackingNumber	OrderQty	ProductID	UnitPrice	Modified Date	
2	43662	20	2E53-4802-85		764	377.51301	7/12/2013 0:00	

If you want to use the sample data, copy the entire table, including the column headings, into an empty Excel file on your computer. For this walkthrough, we named the Excel tabs *Data Table 1* and *DataTable 2*. We named the Excel file *Concatenate_Data.xlsx*.

Data Table 1

Data Table 1

SalesOrder ID	SalesOrderDetailID	TrackingNumber	OrderQty	PI D	SpecialOffer ID	UnitPrice	Modified Date
43659	1	4911-403C-98	1	776	1	1822.4946	7/1/2013 0:00
43659	2	4911-403C-98	3	777	1	1822.4946	7/2/2013 0:00
43659	3	4911-403C-98	1	778	1	1098.9	7/2/2013 0:00
43659	4	4911-403C-98	1	771	1	1835.9946	7/2/2013 0:00
43659	5	4911-403C-98	1	772	1	1835.9946	7/3/2013 0:00
43659	6	4911-403C-98	2	773	1	1835.9946	7/3/2013 0:00
43659	7	4911-403C-98	1	774	1	1835.9946	7/4/2013 0:00
43659	8	4911-403C-98	3	714	1	25.95636	7/4/2013 0:00
43659	9	4911-403C-98	1	716	1	25.95636	7/5/2013 0:00
43659	10	4911-403C-98	6	709	1	5.13	7/5/2013 0:00
43659	11	4911-403C-98	2	712	1	4.66785	7/6/2013 0:00
43659	12	4911-403C-98	4	711	1	18.16785	7/6/2013 0:00

3 Managing data in the app with Data manager

SalesOrder ID	SalesOrderDetail ID	TrackingNumber	OrderQty	PI D	SpecialOffer ID	UnitPrice	Modified Date
43660	13	6431-4D57-83	1	762	1	377.51301	7/7/2013 0:00
43660	14	6431-4D57-83	1	758	1	787.3146	7/7/2013 0:00
43661	15	4E0A-4F89-AE	1	745	1	728.784	7/7/2013 0:00
43661	16	4E0A-4F89-AE	1	743	1	643.23387	7/8/2013 0:00
43661	17	4E0A-4F89-AE	2	747	1	643.23387	7/8/2013 0:00
43661	18	4E0A-4F89-AE	4	712	1	4.66785	7/8/2013 0:00
43661	19	4E0A-4F89-AE	4	715	1	25.95636	7/9/2013 0:00

DataTable 2

Data Table 2

SalesOrderID	SalesOrderDetailID	TrackingNumber	OrderQty	ProductID	UnitPrice	Modified Date
43662	20	2E53-4802-85	3	764	377.51301	7/12/2013 0:00
43662	21	2E53-4802-85	5	770	377.51301	7/12/2013 0:00
43662	22	2E53-4802-85	2	730	165.54438	7/13/2013 0:00
43662	23	2E53-4802-85	4	754	787.3146	7/14/2013 0:00
43662	24	2E53-4802-85	3	725	165.54438	7/14/2013 0:00
43662	25	2E53-4802-85	5	762	377.51301	7/14/2013 0:00
43662	26	2E53-4802-85	3	765	377.51301	7/14/2013 0:00
43662	27	2E53-4802-85	2	768	377.51301	7/15/2013 0:00

3 Managing data in the app with Data manager

SalesOrderID	SalesOrderDetailID	TrackingNumber	OrderQty	ProductID	UnitPrice	Modified Date
43662	28	2E53-4802-85	1	753	1932.2658	7/15/2013 0:00
43663	29	2E53-4802-85	1	756	787.3146	7/16/2013 0:00
43663	30	2E53-4802-85	3	763	377.51301	7/17/2013 0:00
43664	31	2E53-4802-85	1	732	321.2082	7/18/2013 0:00
43664	32	2E53-4802-85	6	758	787.3146	7/19/2013 0:00
43665	33	2E53-4802-85	1	729	165.54438	7/19/2013 0:00

Add data tables to an app


Do the following:


1. Start Qlik Sense.
2. Click **Create new app** in your work area. The **Create new app** window opens.
3. Name your app, and then click **Create**. The app is created. We named our app *ConcatenateExample*
4. Click **Open app**. The app opens and displays a dialog where you can add data.
5. Drag and drop your **Excel** file into the **Add data from files and other sources** dialogue. Your tables are displayed in the **Associations** view of the **Data manager**. Click a bubble to see the data for that table.



*If you add data instead from **Data manager**, you will first be asked to select table fields before being taken to the **Associations** view of the **Data manager**. In this case, select all the*

3 Managing data in the app with Data manager

 *fields for both tables.*



Data Table 1 Concatenate_Data.xlsx Pending add Fields: 8

Data Table 1.SalesOrderID	Data Table 1.SalesOrderDetailID	Data Table 1.TrackingNumber	Data Table 1.OrderQty	PID	SpecialOfferID	Data Table 1.UnitPrice	Data Table 1.Modified Date
43659	1	4911-403C-98	1	776	1	1822.4946	7/1/2013 12:00:00 AM
43659	2	4911-403C-98	3	777	1	1822.4946	7/2/2013 12:00:00 AM
43659	3	4911-403C-98	1	778	1	1098.9	7/2/2013 12:00:00 AM
43659	4	4911-403C-98	1	771	1	1835.9946	7/2/2013 12:00:00 AM
43659	5	4911-403C-98	1	772	1	1835.9946	7/3/2013 12:00:00 AM

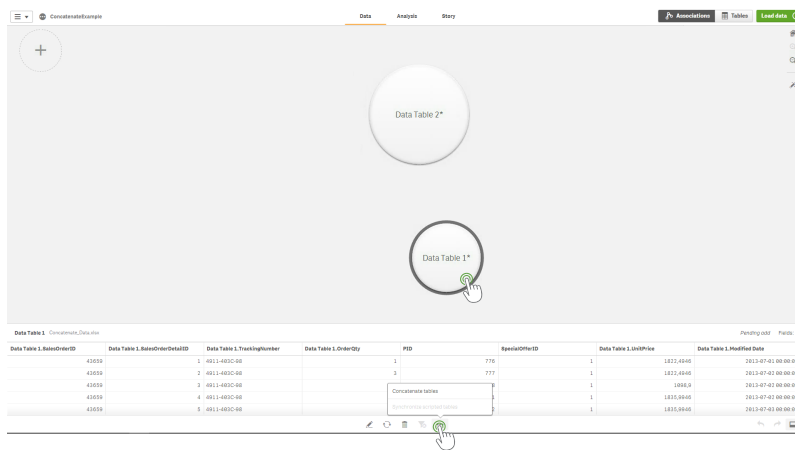
Concatenate tables and load data tables into an app

After the data tables have been added to the app, the tables can be concatenated.

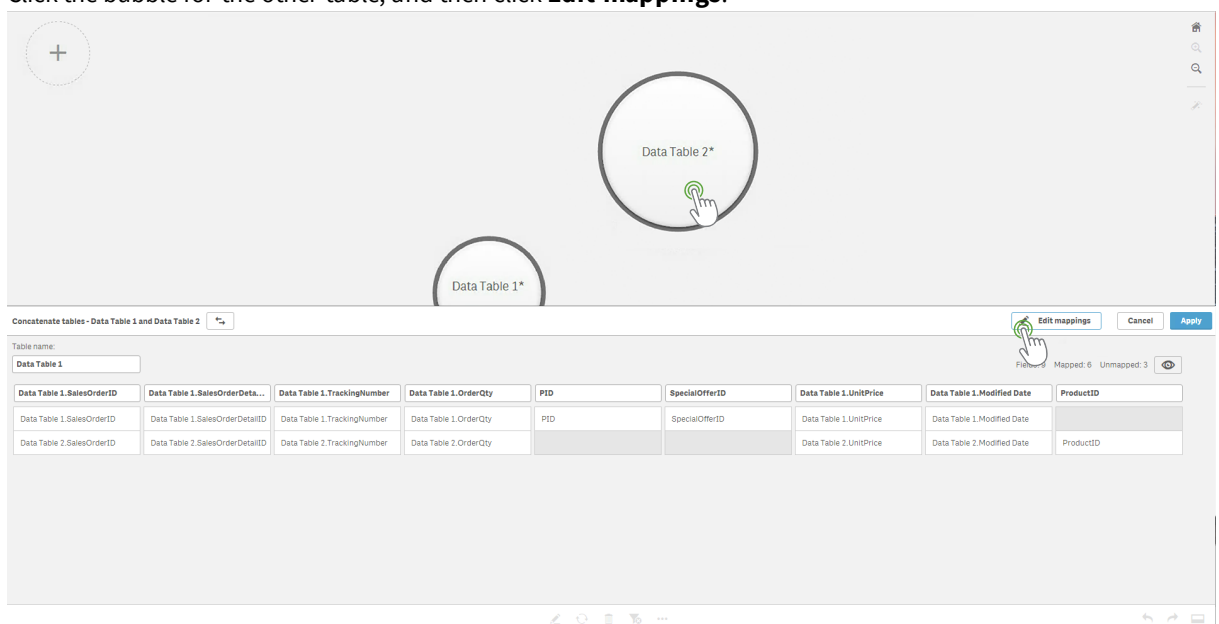
Do the following:


1. In the **Associations** view of **Data manager**, select one table by clicking the bubble. Click **...** and then select **Concatenate tables**.

3 Managing data in the app with Data manager

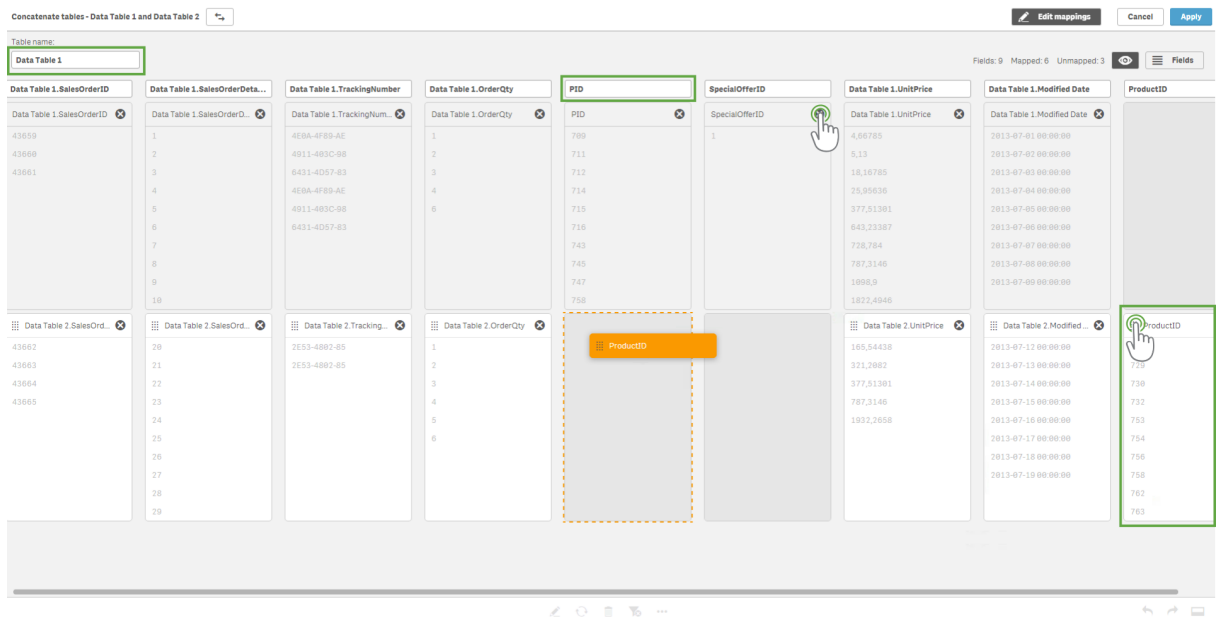


2. Click the bubble for the other table, and then click **Edit mappings**.



3. You can now do the following as required:
 - i. In **Table name**, rename the table that will be created when you combine the tables.
 - ii. Combine fields by dragging and dropping field labels.
 - iii. Rename fields.
 - iv. Delete fields by clicking  for the field.

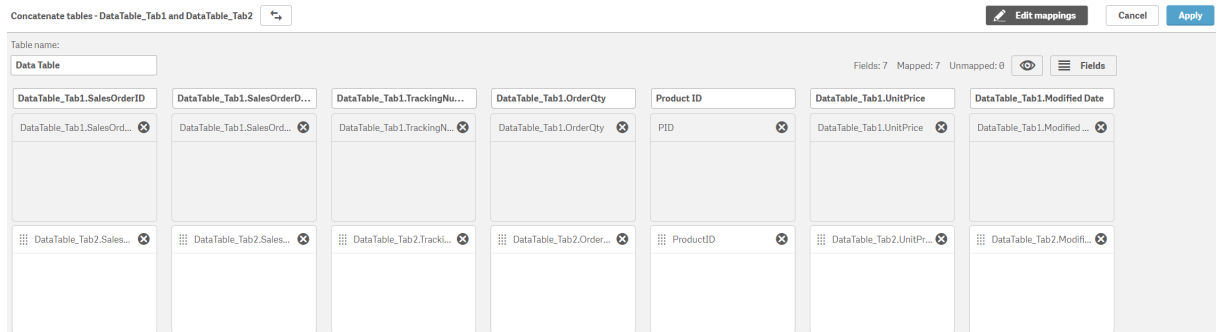
3 Managing data in the app with Data manager



In our example, we did the following:

- i. Renamed our table to Data Table
- ii. Dragged the ProductID label and field to below the PID field, to combine the fields.
- iii. Rename the PID field to Product ID.
- iv. Deleted the SpecialOfferID field.

Our table now looks like this:



4. Click **Apply**. The tables are concatenated on the mapped fields. The * indicates that the data has not yet been loaded into the app.



5. Click **Load data**. A message is displayed indicating that the data was loaded successfully. Click **Edit sheet** to create visualizations using the data.

A step further - adding a new table and concatenating the data fields

The sample data provided above was pasted into two tabs in the same Excel file. However, the tables do not need to be in same file when you want to concatenate fields. The tables can be in separate files that are added to the app. Another table can also be added later, for example if you initially add a table from June, and then later want to add a second table from July.

In this example, we add another table with similar fields to the concatenated table we created above.

Here is the sample data. We named the tab that contains the table *DataTable_Newest*. We named the data file *Concatenate_Data2.xlsx*.

DataTable_Newest

DataTable_Newest

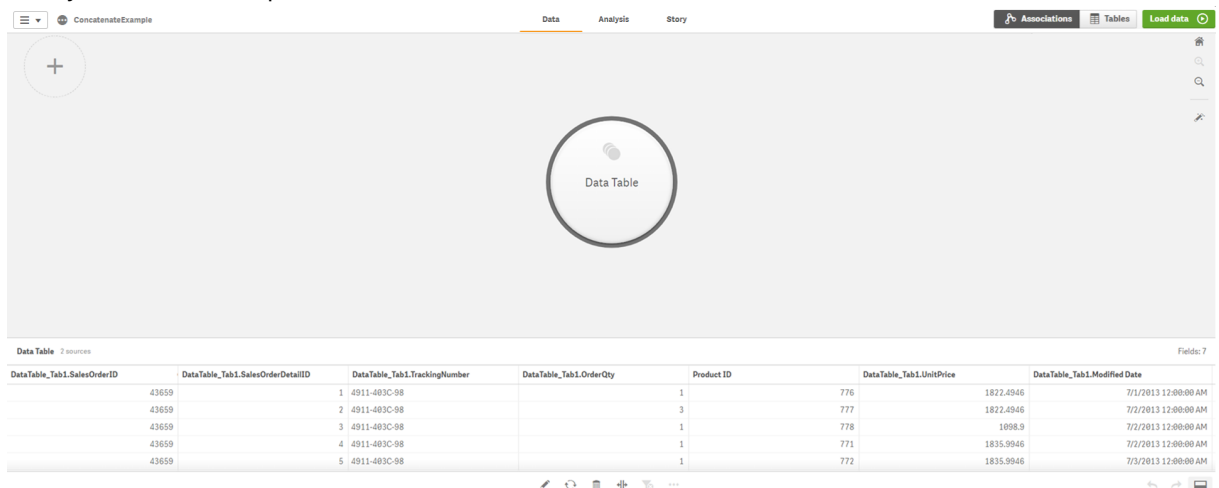
SalesOrderID	SalesOrderDetailID	TrackingNumber	ZIP	OrderQty	ID	UnitPrice	Modified Date
43666	34	568E-472E-9C	20012	3	764	377.51301	7/12/2013 0:00
43666	34	568E-472E-9C	23220	5	770	377.51301	7/12/2013 0:00
43667	35	AB6C-4FF9-9D	30004	2	730	165.54438	7/13/2013 0:00
43668	36	AB6C-4FF9-9D	11215	4	754	787.3146	7/14/2013 0:00
43668	36	AB6C-4FF9-9D	55401	3	725	165.54438	7/14/2013 0:00

3 Managing data in the app with Data manager


SalesOrderID	SalesOrderDetailID	TrackingNumber	ZIP	OrderQty	ID	UnitPrice	Modified Date
43668	36	AB6C-4FF9-9D	20003	5	762	377.51301	7/14/2013 0:00
43668	36	AB6C-4FF9-9D	15213	3	765	377.51301	7/14/2013 0:00
43669	37	C618-4998-BE	33125	2	768	377.51301	7/15/2013 0:00
43669	37	C618-4998-BE	11215	1	753	1932.2658	7/15/2013 0:00

Do the following:

1. From the Qlik Sense hub, click the app that you created in the procedures above. The app opens.
2. Select **Data manager** from the drop-down list in the top toolbar. The **Data manager** opens and the table you created in the procedure above is shown.




3. Click the **+** button to add data.
4. Add the new Excel file to the app by dragging it into the **Attach files to this app** dialogue.



The **Add data** window opens.

5. Click **Add data** to add the data table to the app.
The new table is added to your app.

3 Managing data in the app with Data manager



The screenshot shows the Qlik Data Manager interface. At the top, there are tabs for 'Data', 'Analysis', and 'Story'. Below the tabs, there are two circular nodes representing data tables: 'DataTable_Newest*' and 'Data Table'. Below the nodes, there is a table with 7 columns and 5 rows of data. The columns are: 'DataTable_Tab1.SalesOrderID', 'DataTable_Tab1.SalesOrderDetailID', 'DataTable_Tab1.TrackingNumber', 'DataTable_Tab1.OrderQty', 'Product ID', 'DataTable_Tab1.UnitPrice', and 'DataTable_Tab1.Modified Date'. The data rows are as follows:

DataTable_Tab1.SalesOrderID	DataTable_Tab1.SalesOrderDetailID	DataTable_Tab1.TrackingNumber	DataTable_Tab1.OrderQty	Product ID	DataTable_Tab1.UnitPrice	DataTable_Tab1.Modified Date
43659	1	4911-493C-98		1	776	1822.4946
43659	2	4911-493C-98		3	777	1822.4946
43659	3	4911-493C-98		1	778	1898.9
43659	4	4911-493C-98		1	771	1835.9946
43659	5	4911-493C-98		1	772	1835.9946

6. You can now concatenate the tables, edit the mappings, and load the data.

For more information, see *Concatenate tables and load data tables into an app* (page 89).

3.16 Synchronizing scripted tables in Data manager

By default, scripted tables added in the data load editor cannot use the tools available in **Data manager**. For example, you cannot associate scripted tables with other tables in **Data manager** or transform fields in scripted tables using the data profiling cards. If you synchronize your scripted tables in **Data manager**, you can replace your scripted tables in **Data manager** with managed scripted tables. These tables have access to all the same tools as tables added in **Data manager**, including:

- Editing tables, such as adding calculated fields.
- Transforming fields, such as changing the field types or transforming fields with data profiling cards.
- Transforming tables, such as unpivoting or concatenating tables.

Synchronization and managed scripted tables have the following limitations:

- Scripted tables must be located before the **Auto-generated section** in the data load script to be synchronized as managed scripted tables. Tables after the **Auto-generated section** in the data load script will not be synchronized.
- You cannot use **Select data from source** to change the selection of fields in a managed scripted table.



*Do not synchronize your scripted tables if your data load script contains an **Exit** statement or dynamic fields.*

To convert your scripted tables into managed scripted tables, synchronize your scripted tables in **Data manager**. Synchronization does the following:

- Replaces all synchronized scripted tables as managed scripted tables.
- Deletes any managed scripted tables whose scripted tables have been removed in the data load script.
- Updates any managed scripted tables whose source tables were changed in the data load script.



*If you have synchronized tables, you should not make changes in the data load editor with **Data manager** open in another tab.*



Avoid changing the data load script for tables already synchronized in **Data manager**. If you remove or modify fields in data load editor, you must delete or redo any derived fields or associations in the synchronized table. Derived fields using a removed or modified field, such as a calculated field or fields created by the **Split** card, display null values.

After synchronization, you can use the managed scripted tables in **Data manager** like any other table. **Data manager** prompts you to synchronize again if it detects differences between a managed scripted table and the source scripted table.

To change managed scripted tables back into scripted tables, delete them in **Data manager**. You must repeat the deletion if you synchronize again.

Synchronizing scripted tables


Do the following:

1. In **Data manager**, click **•••**.
Alternatively, select a scripted table.
2. Click **Synchronize scripted tables**.

Managed scripted tables replace all the scripted tables in **Data manager**.

Removing managed scripted tables

Do the following:

1. In **Data manager**, select **Tables** view.
2. On the managed scripted table you want to remove, click .
3. Click **Load data**.

The managed scripted table changes back to a scripted table.

3.17 Managing data associations

Qlik Sense can profile your data to help you create associations between tables. You can make associations by choosing from Qlik Sense's analysis-based recommendations, or you can create your own.

If you want to associate your data, we recommend that you use the **Add data** option with data profiling enabled. This is the default option. You can verify this setting by clicking **•••** beside the **Add data** button in the lower right corner of the Add Data page.

3 Managing data in the app with Data manager

In the **Associations** view of the **Data manager**, your data is illustrated using bubbles, with each bubble representing a data table. The size of the bubble represents the amount of data in the table. The links between the bubbles represent the associations between tables. If there is an association between two tables, you can click the button in the link to view or edit the association.

SKEY	Cost	Sales.Custo...	Date	GrossSales	Invoice Date	Invoice Num...	Item Desc	Sales.Item N...	Margin	Order Number	Promised De...	Sales	Sales Qty
1	-513.15	10012715	12/31/2012	-573.3835	12/31/2012	118960	Cutting Edge Sliced Ham	10096	-60.23	215785	12/31/2012	-573.38	
2	-105.93	10012715	12/31/2012	-204.6638	12/31/2012	118960	Washington Cranberry Juice	10009	-98.73	215785	12/31/2012	-204.66	
3	-88.07	10012715	12/31/2012	-165.8016	12/31/2012	118960	Moms Sliced Ham	10385	-77.73	215785	12/31/2012	-165.8	
4	-43.12	10012715	12/31/2012	-118.3703	12/31/2012	118960	Tip Top Lox	10215	-75.25	215785	12/31/2012	-118.37	
5	-37.98	10012715	12/31/2012	-102.3319	12/31/2012	118960	Just Right Beef Soup	10965	-64.35	215785	12/31/2012	-102.33	
6	-49.37	10012715	12/31/2012	-85.5766	12/31/2012	118960	Fantastic Pumpernickel Bread	10901	-36.21	215785	12/31/2012	-85.58	



*In most cases it is easier to edit table associations in the model view, but you can also edit a single table's associations using the **Associate** option in table edit view.*

For more information, see [Associating data in the table editor \(page 50\)](#).

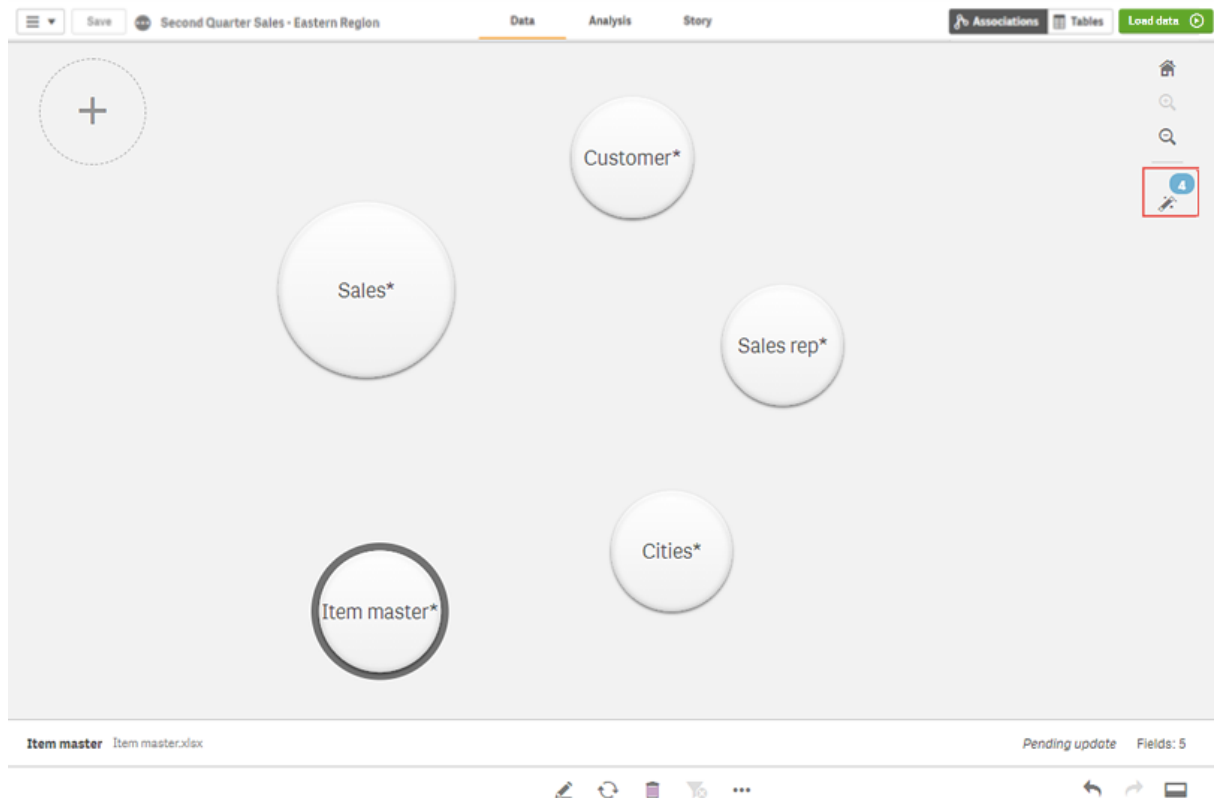
Associating tables using the Recommended associations panel

In many cases, Qlik Sense will recommend associations between data tables. The **Recommended associations** panel lets you view and apply these recommendations.


The **Recommended associations** panel will open by default if any tables are present. It can be closed by clicking on **x** in the upper right corner, or clicking on the . It can be re-opened by clicking on the .

If the panel is closed and recommendations exist, you will see a badge on top of the showing the number of recommendations.

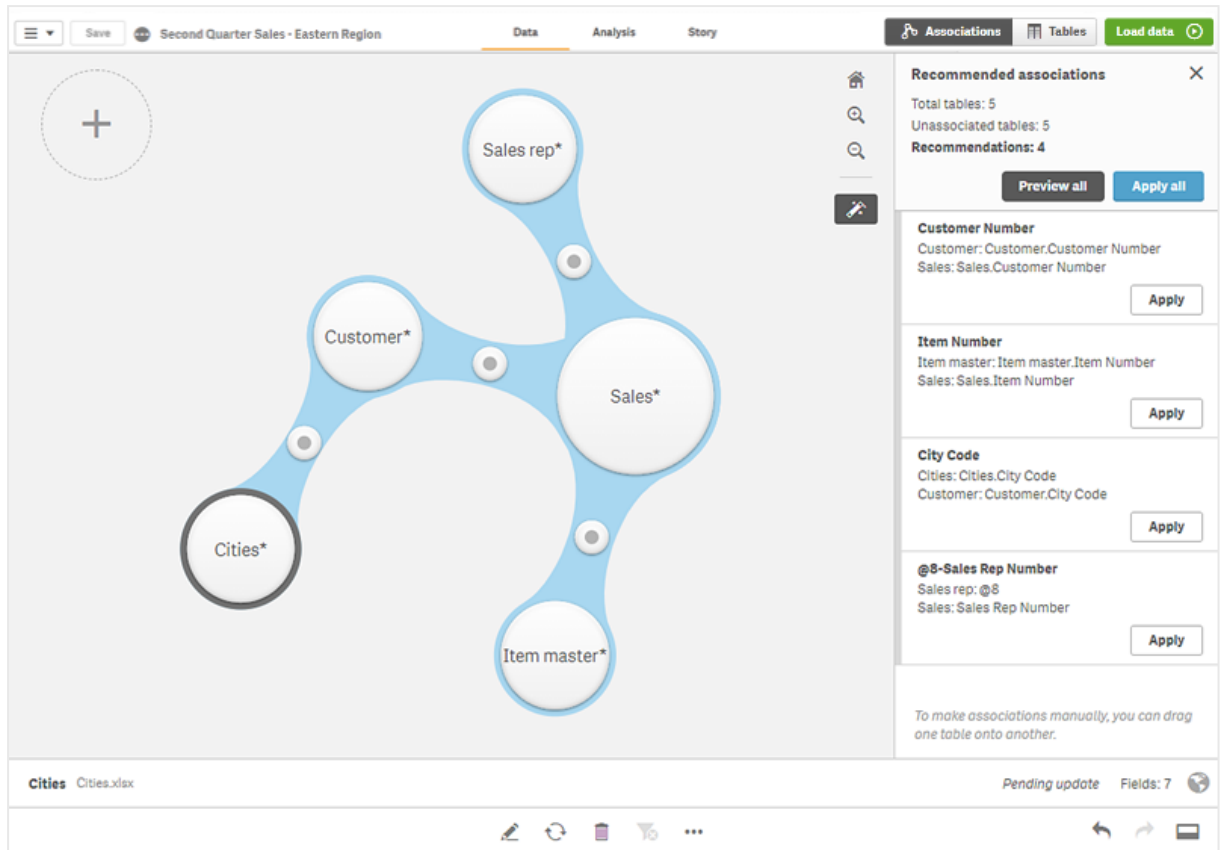
3 Managing data in the app with Data manager




Do the following:

1. Click  in the upper right corner of the associations view, if the **Recommended associations** panel is closed.
The panel appears on the right.
2. You will see the following information:
 - **Total tables:** the total number of tables.
 - **Unassociated tables:** the total number of tables that have no associations.
 - **Recommendations:** the total number of recommended associations.
 - Recommended association details: showing the name of the recommended association, and then table and field names separated by colons
3. Click on a single recommendation to preview it in dark blue.
4. To accept only some of the recommendations, click the **Apply** button for the specific recommendation you need.
5. Click **Preview all** to see how all the recommended associations will affect your data tables. Associations being previewed are highlighted in light blue.

3 Managing data in the app with Data manager



6. Click **Apply all** to apply every recommended association. Associations that have been accepted are highlighted in light grey.
You can click  at the bottom of the screen to see how your tables have changed.

You can now start making visualizations with your data.

Associating tables manually

You can associate tables manually, by dragging them together. When you drag table bubbles towards each other, they will be marked with a green, orange, or red stripe.

- Green: the **Data manager** is very confident about which fields to associate. For example, if two tables have fields labeled "Sales Region", the **Data manager** assumes they should be associated.
- Orange: the **Data manager** is fairly confident that these tables can be associated. For example, if two different fields have different labels, but contain single digit data, the **Data manager** will flag them as orange, because the data types are similar.
- Red: the **Data manager** does not know how to associate these tables. You will have to choose which tables and fields go together in the **Associate tables** editor.

I want to manually associate tables that are green or orange

Do the following:

3 Managing data in the app with Data manager

1. Drag a table to one of the tables marked with green or orange.
2. The association is automatically applied.

The tables are now associated using the recommended fields.

I want to manually associate a tables that are red

You can associate the tables by creating a custom association.

Do the following:

1. Drag a table to one of the tables marked with red.
The **Associate tables** editor opens.
2. In the left table, select which fields to use in the association.
You choose a single field or multiple fields. You can also add delimiter characters to make it easier to interpret the data, or to match a field that already exists. You can see what the association data looks like in the preview.
3. In the right table, select which fields to use to match the selections you made in the left table.
4. Enter a name for the key field that will be created in **Name**.
This new field name cannot be the same as an existing field name in either table.
5. Click **Associate**.

The tables are now associated using your custom association.

Breaking associations

There are two ways of breaking associations that are not a good fit for your data model.

Do the following:


- Click one of the associated tables, and drag it away from the other table until the association breaks. Or you can:
- Click on the link between the two bubbles, and then click the **Delete** button in the bottom panel.

The two tables are no longer associated.


Editing associations

You can edit an existing association between two tables if you need to adjust the data model.

Do the following:

1. Click the circle between the associated tables to open the data panel.
The panel opens with a preview of data in the associated fields.
2. Click .
You will see one or more buttons, each marked with green, orange, or red. Green means the **Data manager** is very confident in the association, orange means somewhat confident, and red means unsure. The current association is marked with grey.

3. Click one of the association buttons:


- Click a recommended association to select it.
- Click an existing custom association () to edit which fields to use in the association.
- Click **Custom association** to create a new association.
This button is only available if there is a recommended association for the table pair.
Custom associations can contain a single field or multiple fields.

You have now changed the association between the table pair.

Previewing data

You can preview tables in the associations view to get a better understanding of the data.

Do the following:

1. Select a table.
2. Click  at the bottom of the view.

The preview pane is displayed with the table data.

Synthetic keys

When two or more data tables have two or more fields in common, this suggests a composite key relationship. Qlik Sense handles this by creating synthetic keys automatically. These keys are anonymous fields that represent all occurring combinations of the composite key.

For more information, see *Synthetic keys (page 133)*.

If adding a table results in any of the following cases, you can only add data with profiling enabled:

- A synthetic key containing more than five fields is created.
- More than ten synthetic keys are created.
- Nested synthetic keys are created, that is, synthetic keys containing other synthetic keys.

These cases indicate that you need to adjust the data tables to resolve the issues.

Limitations

There are some cases where association recommendations are not provided, due to the structure of the loaded tables and the data in the tables. In these cases, you need to adjust the associations in the table editor.

- Many-to-many relationships.
- Field pairs with data does not match well in both directions. This may be the case when you have a small table with a few field values that match a field in a large table 100%, while the match in the other direction is significantly smaller.
- Compound key associations.

3 Managing data in the app with Data manager

Additionally, the **Data manager** will only analyze tables that were added with **Add data**. Tables added using the data load script are not included in the association recommendations, unless they have been synchronized into **Data manager**.


For more information, see *Synchronizing scripted tables in Data manager (page 94)*.

Applying changes and reloading data




Changes that you have made in the **Data manager** will not be available in the app until you have reloaded data. When you reload data, the changes are applied and any new data that you have added is loaded from the external data sources. Data that you loaded previously is not reloaded.

You can reload all the data from the external data sources by using the  button in the **Data manager** footer.



The  button reloads all the data for the selected table. It does not reload all the data for all the tables in the app.

If the data in **Data manager** is out of sync with the app data, the **Load data** button is green. In the **Associations** view, all new or updated tables are indicated with *, and deleted tables are a lighter color of gray. In the **Tables** view, all new, updated, or deleted tables are highlighted in blue and display an icon that shows the status of the table:

- Tables marked with **Pending delete**  will be deleted.
- Tables marked with **Pending update**  will be updated with fields that have been added, renamed, or removed, or the table will be renamed.
- Tables marked with **Pending add**  will be added.

Applying changes

Do the following:

- Click **Load data** to load the changes in the app.

The app data is now updated with changes you made in **Data manager**.

4 Loading data with the data load script

This introduction serves as a brief presentation of how you can load data into Qlik Sense using data load scripts.

Qlik Sense uses a data load script, which is managed in the data load editor, to connect to and retrieve data from various data sources. In the script, the fields and tables to load are specified. It is also possible to manipulate the data structure by using script statements and expressions. It is also possible to load data into Qlik Sense using the data manager, but when you want to create, edit and run a data load script you use the data load editor.

During the data load, Qlik Sense identifies common fields from different tables (key fields) to associate the data. The resulting data structure of the data in the app can be monitored in the data model viewer. Changes to the data structure can be achieved by renaming fields to obtain different associations between tables.

After the data has been loaded into Qlik Sense, it is stored in the app. The app is the heart of the program's functionality and it is characterized by the unrestricted manner in which data is associated, its large number of possible dimensions, its speed of analysis and its compact size. The app is held in RAM when it is open.

Analysis in Qlik Sense always happens while the app is not directly connected to its data sources. So, to refresh the data, you need to run the script to reload the data.

4.1 Interaction between Data manager and the data load script

When you add data tables in the **Data manager**, data load script code is generated. You can view the script code in the **Auto-generated section** of the data load editor. You can also choose to unlock and edit the generated script code, but if you do, the data tables will no longer be managed in the **Data manager**.

By default, data tables defined in the load script are not managed in **Data manager**. That is, you can see the tables in the data overview, but you cannot delete or edit the tables in **Data manager**, and association recommendations are not provided for tables loaded with the script. If you synchronize your scripted tables with **Data manager**, however, your scripted tables are added as managed scripted tables to **Data manager**.



*If you have synchronized tables, you should not make changes in the data load editor with **Data manager** open in another tab.*

You can add script sections and develop code that enhances and interacts with the data model created in **Data manager**, but there are some areas where you need to be careful. The script code you write can interfere with the **Data manager** data model, and create problems in some cases, for example:

- Renaming or dropping tables added with **Data manager** in the script.
- Dropping fields from tables added with **Data manager**.
- Concatenation between tables added with **Data manager** and tables loaded in the script.
- Using the **Qualify** statement with fields in tables added with **Data manager**.

4 Loading data with the data load script

- Loading tables added with **Data manager** using **Resident** in the script.
- Adding script code after the generated code section. The resulting changes in the data model are not reflected in **Data manager**.

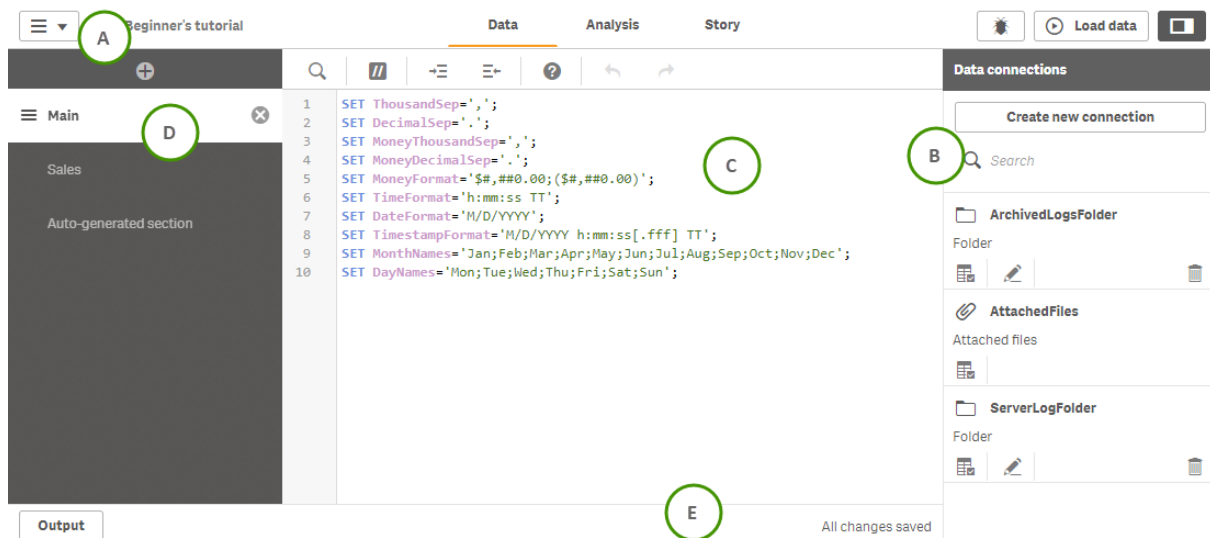
4.2 Using the data load editor

This section describes how to use the data load editor to create or edit a data load script that can be used to load your data model into the app.



The data load script connects an app to a data source and loads data from the data source into the app. When you have loaded the data it is available to the app for analysis. When you want to create, edit and run a data load script you use the data load editor.

You can edit a script manually, or it can be generated by the data manager. If you need to use complex script statements, you need to edit them manually.

Example of the data load editor.



A: Toolbar

Toolbar with the most frequently used commands for the data load editor: global menu,  (debug) and **Load data** . The toolbar also displays the save and data load status of the app.

B: Data connections

Under **Data connections** you can save shortcuts to the data sources (databases or remote files) you commonly use. This is also where you initiate selection of which data to load.

C: Text editor

You can write and edit the script code in the text editor. Each script line is numbered and the script is color coded by syntax components. The text editor toolbar contains commands for **Search and replace**, **Help mode**, **Undo**, and **Redo**. The initial script already contains some pre-defined regional variable settings, for example `SET ThousandSep=,` that you generally do not need to edit.

D: Sections

Divide your script into sections to make it easier to read and maintain. The sections are executed from top to bottom.


If you have added data using **Add data**, you will have a data load script section named **Auto-generated section** that contains the script code required to load the data.

E: Output

Output displays the autosave status and all messages that are generated during script execution.

Quick start

If you want to load a file or tables from a database, you need to complete the following steps in **Data connections**:


1. **Create new connection** linking to the data source (if the data connection does not already exist).
2.  Select data from the connection.

When you have completed the select dialog with **Insert script**, you can select **Load data** to load the data model into your app.




For detailed reference regarding script functions and chart functions, see the Script syntax and chart functions.

Toolbars






The toolbars allow you to perform global actions on your data load script, such as undo/redo, debug, and search/replace. You can also click **Load data**  to reload the data in your app.

Main toolbar

Main toolbar options









Option	Description
	Global menu with navigation options, and actions that you can perform in your app.

4 Loading data with the data load script

Option	Description
Data	<p>Click the tab to perform data tasks. For example you can load data in the Data manager or the Data load editor, or view the data model in the Data model viewer.</p> <p>The Data tab is not available in a published app, unless you are the owner of the app. In that case, you can only open the Data model viewer.</p>
Analysis	Click the tab to perform analysis tasks. For example, you can create or interact with tables and charts.
Story	Click the tab to build stories in your app.
	<p>Show or hide app information, where you can choose to edit app information or open app options and style your app.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p> <i>Once the app has been published, you cannot edit app information nor open app options.</i></p> </div>
	Debug the script.
 Load data	Run the script to load data. The app is automatically saved before reloading.
	Toggle Data connections view.

Editor toolbar

Editor toolbar options

Option	Description
	Search and replace text in the script.
	Comment/uncomment.
	Indent.
	Outdent.
	<p>Activate syntax help mode. In help mode you can click on a syntax keyword (marked in blue) in the editor to access detailed syntax help.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p> <i>It is not possible to edit the script in help mode.</i></p> </div>
	Undo the latest change in the current section (multiple step undo is possible). This is equivalent to pressing Ctrl+Z.
	Redo the latest Undo in the current section. This is equivalent to pressing Ctrl+Y.

Connect to data sources in the data load editor

Data connections in the data load editor provide a way to save shortcuts to the data sources you commonly use: databases, local files or remote files. **Data connections** lists the connections you have saved in alphabetical order. You can use the search/filter box to narrow the list down to connections with a certain name or type.



You can only see data connections that you own, or have been given access rights to read or update. Please contact your Qlik Sense system administrator to acquire access if required.

Creating a new data connection

Do the following:

1. Click **Create new connection**.
2. Select the type of data source you want to create from the drop-down list.
The settings dialog, specific for the type of data source you selected, opens.
3. Enter the data source settings and click **Create** to create the data connection.
The connection name will be appended with your user name and domain to ensure that it is unique.

The data connection is now created with you as the default owner. If you want other users to be able to use the connection in a server installation, you need to edit the access rights of the connection in the Qlik Management Console.




The settings of the connection you created will not be automatically updated if the data source settings are changed. This means you need to be careful about storing user names and passwords, especially if you change settings between integrated Windows security and database logins in the DSN.



*If **Create new connection** is not displayed, it means you do not have access rights to add data connections. Please contact your Qlik Sense system administrator to acquire access if required.*


Deleting a data connection

Do the following:

1. Click  on the data connection you want to delete.
2. Confirm that you want to delete the connection.


The data connection is now deleted.



If  is not displayed, it means you do not have access rights to delete the data connection. Please contact your Qlik Sense system administrator to acquire access if required.

Editing a data connection

Do the following:


1. Click  on the data connection you want to edit.
2. Edit the data connection details. Connection details are specific to the type of connection.

The data connection is now updated.



If you edit the name of a data connection, you also need to edit all existing references (lib://) to the connection in the script, if you want to continue referring to the same connection.



If  is not displayed, it means you do not have access rights to update the data connection. Please contact your Qlik Sense system administrator if required.

Inserting a connect string

Connect strings are required for most connections. Only folder and web file connections do not require connect strings.

Do the following:

- Click  on the connection for which you want to insert a connect string.


A connect string for the selected data connection is inserted at the current position in the data load editor.



You can also insert a connect string by dragging a data connection and dropping it on the position in the script where you want to insert it.

Selecting data from a data connection

If you want to select data from a data connection to load in your app, do the following:

1. **Create new connection** linking to the data source (if the data connection does not already exist).
2.  Select data from the connection.

Referring to a data connection in the script

You can use the data connection to refer to data sources in statements and functions in the script, typically where you want to refer to a file name with a path.

Example: Loading a file from a folder data connection

This example loads the file *orders.csv* from the location defined in the MyData data connection.

Example: Loading a file from a sub-folder

This example loads the file *Customers/cust.txt* from the DataSource data connection folder. Customers is a sub-folder in the location defined in the MyData data connection.

```
LOAD * FROM 'lib://DataSource/Customers/cust.txt';
```

Example: Loading from a web file

This example loads a table from the PublicData web file data connection, which contains the link to the actual URL.

```
LOAD * FROM 'lib://PublicData' (html, table is @1);
```

Example: Loading from a database

This example loads the table Sales_data from the MyDataSource database connection.

```
LIB CONNECT TO 'MyDataSource';  
LOAD *;  
SQL SELECT * FROM `Sales_data`;
```

Where is the data connection stored?

Connections are stored using the Qlik Sense Repository Service. You can manage data connections with the Qlik Management Console in a Qlik Sense server deployment. The Qlik Management Console allows you to delete data connections, set access rights and perform other system administration tasks.



In Qlik Sense Desktop, all connections are saved in the app without encryption. This includes possible details about user name, password, and file path that you have entered when creating the connection. This means that all these details may be available in plain text if you share the app with another user. You need to consider this while designing an app for sharing.


Select data in the data load editor

You can select which fields to load from files or database tables and which views of the data source you want by using **Select data** in the data load editor. As well as selecting fields, you can also rename fields in the dialog. When you have finished selecting fields, you can insert the generated script code into your script.

Some data sources, such as a CSV file, contain a single table, while other data sources, such as Microsoft Excel spreadsheets or databases can contain several tables.




*Do not add a table in **Data load editor** that has already been added as a scripted table with the same name and same columns in **Data manager**.*

You open **Select data** by clicking  on a data connection in the data load editor.

Selecting data from a database

When selecting data from a database, the data source can contain several tables.

Do the following:

1. Click  on a database connection in the data load editor.
The select data dialog is displayed.
2. Select a **Database** from the drop-down list.
Some selection dialogs do not have a **Database** drop-down list because the database name is entered when the connection is configured.
3. Select **Owner** of the database.
The list of **Tables** is populated with views and tables available in the selected database.
Some databases do not require that owners be specified in the data selection process.
4. Select a table.
5. Select the fields you want to load by checking the box next to each field you want to load.
You can select all fields in the table by checking the box next to the table name.



You can edit the field name by clicking on the existing field name and typing a new name. This may affect how the table is linked to other tables, as they are joined on common fields by default.

6. Select additional tables if you want to add data from them.



You cannot rename fields in the data selection wizard at the same time as you filter for fields by searching. You have to erase the search string in the text box first.




It is not possible to rename two fields in the same table so that they have identical names.

7. When you have finished your data selection, do the following:
 - Click **Insert script**.
The data selection window is closed, and the LOAD /SELECT statements are inserted in the script in accordance with your selections.

Selecting data from a Microsoft Excel spreadsheet

When you select data from a Microsoft Excel spreadsheet, the file can contain several sheets. Each sheet is loaded as a separate table. An exception is if the sheet has the same field/column structure as another sheet or loaded table, in which case, the tables are concatenated.

Do the following:

1. Click  on the appropriate folder connection in the data load editor.
The select file dialog is displayed.

4 Loading data with the data load script

2. Select a file from the list of files accessible to this folder connection.
3. Select the first sheet to select data from. You can select all fields in a sheet by checking the box next to the sheet name.
4. Make sure you have the appropriate settings for the sheet:

Settings to assist you with interpreting the table data correctly

UI item	Description
Field names	Set to specify if the table contains Embedded field names or No field names . Typically in an Excel spreadsheet, the first row contains the embedded field names. If you select No field names , fields will be named A,B,C...
Header size	Set to the number of rows to omit as table header, typically rows that contain general information that is not in a columnar format.

Example

My spreadsheet looks like this:

Spreadsheet

Machine:	AEJ12B	-	-
Date:	2015-10-05 09	-	-
Timestamp	Order	Operator	Yield
2015-10-05 09:22	00122344	A	52
2015-10-05 10:31	00153534	A	67
2015-10-05 13:46	00747899	B	86

In this case you probably want to ignore the two first lines, and load a table with the fields Timestamp, Order, Operator, and Yield. To achieve this, use these settings:

Settings to ignore the two first lines and load the fields

UI item	Description
Header size	2 This means that the first two lines are considered header data and ignored when loading the file. In this case, the two lines starting with Machine: and Date: are ignored, as they are not part of the table data.
Field names	Embedded field names. This means that the first line that is read is used as field names for the respective columns. In this case, the first line that will be read is the third line because the two first lines are header data.

5. Select the fields you want to load by checking the box next to each field you want to load.



You can edit the field name by clicking on the existing field name and typing a new name. This may affect how the table is linked to other tables, as they are joined by common fields by default.

6. When you are done with your data selection, do the following:

- Click **Insert script**.

The data selection window is closed, and the LOAD /SELECT statements are inserted in the script in accordance with your selections.




*You can also use a Microsoft Excel file as data source using the ODBC interface. In that case you need to use an **ODBC** data connection instead of a **All files** data connection.*

Selecting data from a table file

You can select data from a large number of data files:

- Text files, where data in fields is separated by delimiters such as commas, tabs or semicolons (comma-separated variable (CSV) files).
- HTML tables.
- XML files.
- KML files.
- Qlik native QVD and QVX files.
- Fixed record length files.
- DIF files (Data Interchange Format).

Do the following:

1. Click  on the appropriate folder connection in the data load editor. The select file dialog is displayed.
2. Select a file from the list of files accessible to this folder connection.
3. Make sure that the appropriate file type is selected in **File format**.
4. Make sure you have the appropriate settings for the file. File settings are different for different file types.
5. Select the fields you want to load by checking the box next to each field you want to load. You can also select all fields in a file by checking the box next to the sheet name.



You can edit the field name by clicking on the existing field name and typing a new name. This may affect how the table is linked to other tables, as they are joined by common fields by default.

6. When you are done with your data selection, do the following:

- Click **Insert script**.

The data selection window is closed, and the LOAD /SELECT statements are inserted in the script in accordance with your selections.

Choosing settings for file types

Delimited table files

These settings are validated for delimited table files, containing a single table where each record is separated by a line feed, and each field is separated with a delimited character, for example a CSV file.

File format settings for delimited table files

UI item	Description
File format for delimited table files	Set to Delimited or Fixed record . When you make a selection, the select data dialog will adapt to the file format you selected.
Field names	Set to specify if the table contains Embedded field names or No field names .
Delimiter	Set the Delimiter character used in your table file.
Quoting	Set to specify how to handle quotes: None = quote characters are not accepted Standard = standard quoting (quotes can be used as first and last characters of a field value) MSQ = modern-style quoting (allowing multi-line content in fields)
Header size	Set the number of lines to omit as table header.
Character set	Set character set used in the table file.
Comment	Data files can contain comments between records, denoted by starting a line with one or more special characters, for example //. Specify one or more characters to denote a comment line. Qlik Sense does not load lines starting with the character(s) specified here.
Ignore EOF	Select Ignore EOF if your data contains end-of-file characters as part of the field value.

Fixed record data files

Fixed record data files contain a single table in which each record (row of data) contains a number of columns with a fixed field size, usually padded with spaces or tab characters.

You can set the field break positions in two different ways:

- Manually, enter the field break positions separated by commas in **Field break positions**. Each position marks the start of a field.

Example: 1,12,24

- Enable **Field breaks** to edit field break positions interactively in the field data preview. **Field break positions** is updated with the selected positions. You can:
 - Click in the field data preview to insert a field break.
 - Click on a field break to delete it.
 - Drag a field break to move it.

File format settings for fixed record data files

UI item	Description
Field names	Set to specify if the table contains Embedded field names or No field names .
Header size	Set Header size to the number of lines to omit as table header.
Character set	Set to the character set used in the table file.
Tab size	Set to the number of spaces that one tab character represents in the table file.
Record line size	Set to the number of lines that one record spans in the table file. Default is 1.

HTML files

HTML files can contain several tables. Qlik Sense interprets all elements with a <TABLE> tag as a table.

File format settings for HTML files

UI item	Description
Field names	Set to specify if the table contains Embedded field names or No field names .
Character set	Set the character set used in the table file.

XML files

You can load data that is stored in XML format.

There are no specific file format settings for XML files.

QVD files

You can load data that is stored in QVD format. QVD is a native Qlik format and can only be written to and read by Qlik Sense or QlikView. The file format is optimized for speed when reading data from a Qlik Sense script but it is still very compact.

There are no specific file format settings for QVD files.

QVX files

You can load data that is stored in Qlik data eXchange (QVX) format. QVX files are created by custom connectors developed with the Qlik QVX SDK.

There are no specific file format settings for QVX files.

KML files

You can load map files that are stored in KML format, to use in map visualizations.

There are no specific file format settings for KML files.

Previewing scripts

The statements that will be inserted are displayed in the script preview, which you can choose to hide by clicking **Preview script**.

Including LOAD statements

If **Include LOAD statement** is selected, SELECT statements are generated with preceding LOAD statements using the SELECT statements as input.



If you rename fields in a table, a LOAD statement will be inserted automatically regardless of this setting.

Edit the data load script

You write the script in the text editor of the **Data load editor**. Here you can make manual changes to the **LOAD** or **SELECT** statements you have generated when selecting data, and type new script.

The script, which must be written using the Qlik Sense script syntax, is color coded to make it easy to distinguish the different elements. Comments are highlighted in green, whereas Qlik Sense syntax keywords are highlighted in blue. Each script line is numbered.


There are a number of functions available in the editor to assist you in developing the load script, and they are described in this section.

Accessing syntax help for commands and functions

There are several ways to access syntax help for a Qlik Sense syntax keyword.

Accessing the help portal

You can access detailed help in the Qlik Sense help portal in two different ways.

- Click  in the toolbar to enter syntax help mode. In syntax help mode you can click on a syntax keyword (marked in blue and underlined) to access syntax help.
- Place the cursor inside or at the end of the keyword and press Ctrl+H.



You cannot edit the script in syntax help mode.

Using the auto-complete function

If you start to type a Qlik Sense script keyword, you get an auto-complete list of matching keywords to select from. The list is narrowed down as you continue to type, and you can select from templates with suggested syntax and parameters. A tooltip displays the syntax of the function, including parameters and additional

statements, as well as a link to the help portal description of the statement or function.



You can also use the keyboard shortcut Ctrl+Space to show the keyword list, and Ctrl+Shift+Space to show a tooltip.

Inserting a prepared test script

You can insert a prepared test script that will load a set of inline data fields. You can use this to quickly create a data set for test purposes.

Do the following:



- Press Ctrl + 00.

The test script code is inserted into the script.

Indenting code

You can indent the code to increase readability.

Do the following:

1. Select one or several lines to change indentation.
2. Click  to indent the text (increase indentation) or, click  to outdent the text (decrease indentation).



You can also use keyboard shortcuts:

Tab (indent)





Shift+Tab (outdent)

Searching and replacing text

You can search and replace text throughout script sections.

Searching for text

Open the data load editor. Do the following:





1. Click  in the toolbar.
The search drop-down dialog is displayed.
2. In the search box, type the text you want to find.
The search results are highlighted in the current section of the script code. Also, the number of text instances found is indicated next to the section label.
3. You can browse through the results by clicking  and .
4. Click  in the toolbar to close the search dialog.



Also, you can select **Search in all sections** to search in all script sections. The number of text instances found is indicated next to each section label. You can select **Match case** to make case sensitive searches.

Replacing text

Do the following:

1. Click  in the toolbar.
The search drop-down dialog is displayed.
2. Type the text you want to find in the search box.
3. Type the replacement text in the replace box and click **Replace**.
4. Click  to find next instance of search text and do one of the following:
 - Click **Replace** to replace text.
 - Click  to find next.
5. Click  in the toolbar to close the search dialog.



You can also click **Replace all in section** to replace all instances of the search text in the current script section. The replace function is case sensitive, and replaced text will have the case given in the replace field. A message is shown with information about how many instances that were replaced.


Commenting in the script

You can insert comments in the script code, or deactivate parts of the script code by using comment marks. All text on a line that follows to the right of // (two forward slashes) will be considered a comment and will not be executed when the script is run.

The data load editor toolbar contains a shortcut for commenting or uncommenting code. The function works as a toggle. That is, if the selected code is commented out it will be commented, and vice versa.

Commenting

Do the following:


1. Select one or more lines of code that are not commented out, or place the cursor at the beginning of a line.
2. Click  , or press Ctrl + K.

The selected code is now commented out.

Uncommenting

Do the following:

1. Select one or more lines of code that are commented out, or place the cursor at the beginning of a commented line.

2. Click , or press Ctrl + K.

The selected code will now be executed with the rest of the script.



There are more ways to insert comments in the script code:

- Using the **Rem** statement.
- Enclosing a section of code with `/*` and `*/`.

Example:

```
Rem This is a comment ;
```

```
/* This is a comment  
that spans two lines */
```

```
// This is a comment as well
```

Selecting all code

You can select all code in the current script section.

Do the following:


- Press Ctrl + A.

All script code in the current section is selected.

Organizing the script code

You can divide your script into sections to organize the structure. The script is executed in the order of the sections from top to bottom. A script must contain at least one section.

Sections Main, Dictionary and Calendar in the Data load editor.



Section	Line	Code
Main	1	SET I
	2	SET D
	3	SET M
	4	SET M
	5	SET M
	6	SET I
	7	SET D
	8	SET I
	9	SET M

If you have added data using **Add data**, you will have a data load script section named **Auto-generated section** that contains the script code required to load the data.

Creating a new script section

You can insert new script sections to organize your code.

Do the following:

- Click .

The new section is inserted after the currently selected section.


Deleting a script section

You can delete a script section, including all code in the section.



Deleting a script section cannot be undone.

Do the following:

- Click  next to the section tab to delete it. You need to confirm the deletion.
The section is now deleted.

Renaming a script section

You can rename a script section.

Do the following:


1. Click on the section name and type to edit the name.
2. Press Enter or click outside the section when you are finished.

The section is now renamed.

Rearranging script sections

You can rearrange the order of sections to change the script execution order.

Do the following:

1. Select the section you want to move.
2. Put the cursor on the  drag bars and drag the section to rearrange the order.


The sections have now been rearranged.

Debug the data load script

You can use the debugging utilities in the Data load editor to step through the execution of your script by using breakpoints, which enable you to inspect variable values and output from the script execution. You can select whether you want to view any or all of **Output**, **Variables** and **Breakpoints**.

To show the debug panel, do the following:

4 Loading data with the data load script

- Click  in the data load editor toolbar.
The debug panel is opened at the bottom of the Data load editor.







You can't create connections, edit connections, select data, save the script or load data while you are running in debug mode. Debug mode begins with debug execution and continues until the script is executed or execution has been ended.


Debug toolbar


The Data load editor debug panel contains a toolbar with the following options to control the debug execution:

Debug toolbar options

UI item	Description
Limited load	Enable this to limit how many rows of data to load from each data source. This is useful to reduce execution time if your data sources are large. Enter the number of rows you want to load.  <i>This only applies to physical data sources. Automatically generated and Inline loads will not be limited, for example.</i>
	Start or continue execution in debug mode until the next breakpoint is reached.
	Step to the next line of code.
	End execution here. If you end before all code is executed, the resulting data model will only contain data up to the line of code where execution ended.

Output

Output displays all messages that are generated during debug execution. You can select to lock the output from scrolling when new messages are displayed by clicking .

Additionally, the output menu () contains the following options:

Output menu options

UI item	Description
Clear	Click this to delete all output messages.
Select all text	Click this to select all output messages.
Scroll to bottom	Click this to scroll to the last output message.

Variables

Variables lists all reserved variables, system variables and variables defined in the script, and displays the current values during script execution.

Setting a variable as favorite

If you want to inspect specific variables during execution, you can set them as favorites. Favorite variables are displayed at the top of the variable list, marked by a yellow star. To set a variable as favorite, do the following:

- Click on the ★ next to a variable.
The ★ is now yellow, and the variable moved to the top of the variable list.

Filtering variables

You can apply a filter to show only a selected type of variables by using the following options in the variables menu (☰):

Variables menu options

UI item	Description
Show all variables	Click this to show all types of variables.
Show system variables	Click this to show system variables. System variables are defined by Qlik Sense, but you can change the variable value in the script.
Show reserved variables	Click this to show reserved variables. Reserved variables are defined by Qlik Sense and the value cannot be changed.
Show user defined variables	Click this to show user defined variables. User defined variables are variables that you have defined in the script.

Breakpoints

You can add breakpoints to your script to be able to halt debug execution at certain lines of code and inspect variable values and output messages at this point. When you have reached a breakpoint, you can choose to stop execution, continue until the next breakpoint is reached, or step to the next line of code. All breakpoints in the scripts are listed, with a reference to section and line number.

Adding a breakpoint

To add a breakpoint at a line of code, do one of the following:



- In the script, click in the area directly to the right of the line number where you want to add a breakpoint. A ● next to the line number will indicate that there is a breakpoint at this line.



You can add breakpoints even when the debug panel is closed.


Deleting breakpoints


You can delete a breakpoint by doing either of the following:

- In the script, click on a  next to the line number.
- In the list of breakpoints, click  next to a breakpoint.

You can also click  and select **Delete all** to delete all breakpoints from the script.

Enabling and disabling breakpoints

When you create a breakpoint it is enabled by default, which is indicated by  next to the breakpoint in the list of breakpoints. You can enable and disable individual breakpoints by selecting and deselecting them in the list of breakpoints.


You also have the following options in the breakpoints menu ():


- **Enable all**
- **Disable all**

Saving the load script

When you save a script the entire app is saved, but data is not automatically reloaded.

Data load editor saves your work automatically as you make changes to your load script. You can force a save by pressing CTRL+S.


When the script is saved, the app will still contain old data from the previous reload, which is indicated in the toolbar. If you want to update the app with new data, click **Load data**  in the data load editor toolbar.

When you save a script, it is automatically checked for syntax errors. Syntax errors are highlighted in the code, and all script sections containing syntax errors are indicated with  next to the section label.



The script is automatically saved to the app when data is loaded.

Run the script to load data

Click **Load data**  in the toolbar to run the script and reload data in the app. The app is automatically saved before loading the data.

The **Data load progress** dialog is displayed, and you can **Abort** the load. When the data load has completed, the dialog is updated with status (**Finished successfully** or **Data load failed**) and a summary of possible errors and warnings, such as for synthetic keys. The summary is also displayed in **Output**, if you want to view it after the dialog is closed.



If you want the **Data load progress** dialog to always close automatically after a successful execution, select **Close when successfully finished**.

Keyboard shortcuts in the Data load editor

There are a number of keyboard shortcuts you can use to work effectively and easily in the **Data load editor** environment.

Keyboard shortcuts



Keyboard shortcuts are expressed assuming that you are working in Windows. For Mac OS use Cmd instead of Ctrl.

Keyboard shortcuts in the Data load editor

Shortcut	Action
Ctrl+P	Prints the current view or active sheet/story.
Ctrl+C	Copies the selected item to the clipboard.
Ctrl+X	Cuts the selected item and copies it to the clipboard. When using the Google Chrome browser: if the cursor is put in front of a row in the data load editor or in the expression editor, without selecting anything, the entire row is cut.
Ctrl+V	Pastes the most-recently copied item from the clipboard.
Ctrl+Z	Undo action. Repeat to undo earlier actions.
Ctrl+H	Opens online help in the context of the current selected function, while in the data load editor or the expression editor.
Ctrl+F	Shows the Find search entry field or hides it if it is open.
Ctrl+S	Saves changes.
Enter/Return in search field	Searches for the next instance of the search string.
Enter/Return in replace field	Replaces the selected instance of the search string.
Esc in search or replace field	Closes the Find search entry field.
Ctrl-shift-Enter/Return	Reloads data.
Alt+F11	Expands the script editor to full screen.
Alt+F5	Shows the debug tools or hides them if they are visible.

4 Loading data with the data load script

Shortcut	Action
Alt+F6	Runs debugging if the debug tool is on .
Alt+F7	Proceeds to the next step in the debugger if the debug tool is on.
Alt+F8	Stops the debugger if the debug tool is on.
F9	Toggles insertion and removal of a debug breakpoint.
Alt+F9	Shows the left panel or hides it if it is visible.
Alt+F10	Shows the right panel or hides it if it is visible.
Alt+PgUp	Navigates to the previous section.
Alt+PgDn	Navigates to the next section.
Alt+Insert	Inserts a new section in the script.
Ctrl+K	Toggles insertion and removal of a code comment line.
Tab	Indents a line in the script.
Shift+Tab	Outdents a line in the script.
Alt+1	Shows the Output panel or hides it if it is visible.
Alt+2	Shows the Variables panel or hides it if it is visible, if the debug tool is on .
Alt+3	Shows the Breakpoints panel or hides it if it is visible, if the debug tool is on .
Ctrl+space	Automatically completes with the auto-text string.
Ctrl+Shift+space	Displays a tooltip. (Not supported in Qlik Sense Desktop)

4.3 Understanding script syntax and data structures

Extract, transform and load

In general, the way you load data into the app can be explained by the extract, transform and load process:

- Extract

The first step is to extract data from the data source system. In a script, you use **SELECT** or **LOAD** statements to define this. The differences between these statements are:

- **SELECT** is used to select data from an ODBC data source or OLE DB provider. The **SELECT** SQL statement is evaluated by the data provider, not by Qlik Sense.
- **LOAD** is used to load data from a file, from data defined in the script, from a previously loaded table, from a web page, from the result of a subsequent **SELECT** statement or by generating data automatically.

- Transform

The transformation stage involves manipulating the data using script functions and rules to derive the desired data model structure. Typical operations are:

- Calculating new values
 - Translating coded values
 - Renaming fields
 - Joining tables
 - Aggregating values
 - Pivoting
 - Data validation
- Load

In the final step, you run the script to load the data model you have defined into the app.

Your goal should be to create a data model that enables efficient handling of the data in Qlik Sense. Usually this means that you should aim for a reasonably normalized star schema or snowflake schema without any circular references, that is, a model where each entity is kept in a separate table. In other words a typical data model would look like this:

- a central fact table containing keys to the dimensions and the numbers used to calculate measures (such as number of units, sales amounts, and budget amounts).
- surrounding tables containing the dimensions with all their attributes (such as products, customers, categories, calendar, and suppliers).



In many cases it is possible to solve a task, for example aggregations, either by building a richer data model in the load script, or by performing the aggregations in the chart expressions. As a general rule, you will experience better performance if you keep data transformations in the load script.



It's good practice to sketch out your data model on paper. This will help you by providing structure to what data to extract, and which transformations to perform.

Data loading statements

Data is loaded by **LOAD** or **SELECT** statements. Each of these statements generates an internal table. A table can always be seen as a list of something, each record (row) then being a new instance of the object type and each field (column) being a specific attribute or property of the object.

The differences between these statements are:

- **SELECT** is used to select data from an ODBC data source or OLE DB provider. The **SELECT SQL** statement is evaluated by the data provider, not by Qlik Sense.
- **LOAD** is used to load data from a file, from data defined in the script, from a previously loaded table, from a web page, from the result of a subsequent **SELECT** statement or by generating data automatically.

Rules

The following rules apply when loading data into Qlik Sense:

- Qlik Sense does not make any difference between tables generated by a **LOAD** or a **SELECT** statement. This means that if several tables are loaded, it does not matter whether the tables are loaded by **LOAD** or **SELECT** statements or by a mix of the two.
- The order of the fields in the statement or in the original table in the database is arbitrary to the Qlik Sense logic.
- Field names are used in the further process to identify fields and making associations. These are case sensitive, which often makes it necessary to rename fields in the script.

Execution of the script

For a typical **LOAD** or **SELECT** statement the order of events is roughly as follows:

1. Evaluation of expressions
2. Renaming of fields by **as**
3. Renaming of fields by **alias**
4. Qualification of field names
5. Mapping of data if field name matches
6. Storing data in an internal table

Fields

Fields are the primary data-carrying entities in Qlik Sense. A field typically contains a number of values, called field values. In database terminology we say that the data processed by Qlik Sense comes from data files. A file is composed of several fields where each data entry is a record. The terms file, field and record are equivalent to table, column and row respectively. The Qlik Sense AQL logic works only on the fields and their field values.

Field data is retrieved by script via **LOAD**, **SELECT** or **Binary** statements. The only way of changing data in a field is by re-executing the script. The actual field values can not be manipulated by the user from the layout or by means of automation. Once read into Qlik Sense they can only be viewed and used for logical selections and calculations.

Field values consist of numeric or alphanumeric (text) data. Numeric values actually have dual values, the numeric value and its current, formatted text representation. Only the latter is displayed in sheet objects etc.

The content of a field can be represented in a filter pane.

Derived fields

If you have a group of fields that are related, or if fields carry information that can be broken up into smaller parts that are relevant when creating dimensions or measures, you can create field definitions that can be used to generate derived fields. One example is a date field, from which you can derive several attributes, such as year, month, week number, or day name. All these attributes can be calculated in a dimension expression using Qlik Sense date functions, but an alternative is to create a calendar definition that is common for all fields of date type. Field definitions are stored in the data load script.



Default calendar field definitions for Qlik Sense are included in `autoCalendar` for date fields loaded using **Data manager**. For more information, see *Adding data to the app* (page 15).

Declare the calendar field definitions

You use the **Declare** statement to create a definition of the derived fields. This is where you define the different attributes of the field, in this case date related attributes. Each field is described as `<expression> As field_name tagged tag`. Setting one or more tags is optional, but it can affect the sort order of the derived field. Use `$1` to reference the data field from which the derived fields should be generated.



Do not use `autoCalendar` as name for calendar field definitions, as this name is reserved for auto-generated calendar templates.

Calendar:

```
DECLARE FIELD DEFINITION TAGGED '$date'
```

```
Parameters
```

```
    first_month_of_year = 1
```

```
Fields
```

```
    Year($1) As Year Tagged ('$numeric'),
```

```
    Month($1) as Month Tagged ('$numeric'),
```

```
    Date($1) as Date Tagged ('$date'),
```

```
    week($1) as week Tagged ('$numeric'),
```

```
    weekday($1) as weekday Tagged ('$numeric'),
```

```
    DayNumberOfYear($1, first_month_of_year) as DayNumberOfYear Tagged ('$numeric')
```

```
Groups
```

```
    Year, week, weekday type drilldown as YearWeekDayName,
```

```
    Year, Month, Date type collection as YearMonthDate;
```

Map data fields to the calendar with Derive

The next step is to use the **Derive** statement to map existing data fields to the calendar. This will create the derived fields. You can do this in three alternative ways in the data load script:

- Map specific fields by field name.
`DERIVE FIELDS FROM FIELDS OrderDate,ShippingDate USING Calendar;`
- Map all fields with one or more specific field tags.
`DERIVE FIELDS FROM EXPLICIT TAGS '$date' USING Calendar;`
- Map all fields that are tagged with one of the tags of the field definition (`$date` in the example above).
`DERIVE FIELDS FROM IMPLICIT TAG USING Calendar;`

In this case, you could use any of the three examples here.

Use the derived date fields in a visualization

Qlik Sense is prepared to recognize derived date fields if you have created a calendar definition and mapped the fields like in the example here. They are available in the **Date & time fields** section of the **Fields** asset panel. You will also find all derived fields in the expression editor and when you create or edit dimensions.

Field tags

Field tags provide the possibility of adding metadata to the fields in your data model. There are two different types of field tags:

- System field tags
System field tags are generated automatically when the script is executed and data is loaded. Some of the tags can be manipulated in the script. System tags are always preceded by a \$ sign.
- Custom field tags
You can add custom tags to fields in the data load script, using the **Tag** statement. Custom tags may not use the same name as any system tag.

System field tags

The following system field tags are generated automatically when data is loaded.

System field tags

Tag	Description	Can be manipulated in the script
\$system	System field that is generated by Qlik Sense during script execution.	No
\$key	Key field providing a link between two or more tables.	No
\$keypart	The field is part of one or more synthetic keys.	No
\$syn	Synthetic key	No
\$hidden	Hidden field, that is, it is not displayed in any field selection list when creating visualizations, dimensions or measures. You can still use hidden fields in expressions, but you need to type the field name. You can use the HidePrefix and HideSuffix system variables to set which fields to hide.	Yes
\$numeric	All (non-NULL) values in the field are numeric.	Yes
\$integer	All (non-NULL) values in the field are integers.	Yes
\$text	No values in the field are numeric.	Yes
\$ascii	Field values contain only standard ASCII characters.	Yes
\$date	All (non-NULL) values in the field can be interpreted as dates (integers).	Yes
\$timestamp	All (non-NULL) values in the field can be interpreted as time stamps.	Yes
\$geoname	Field values contain names of geographical locations, related to a point field (\$geopoint) and/or an area field (\$geomultipolygon).	Yes
\$geopoint	Field values contain geometry point data, representing points on a map in the format [longitude, latitude].	Yes
\$geomultipolygon	Field values contain geometry polygon data, representing areas on a map.	Yes

Derived field tags

The following tags can be used when you declare derived fields to specify how to use and display the fields on a contiguous axis in a line chart. You can manipulate the tags in the data load script.

Derived field tags

Tag	Description
\$axis	The \$axis tag is used to specify that the field should generate a tick on the contiguous axis of the chart.
\$qualified \$simplified	<p>You can specify a qualified and a simplified version of an axis label by deriving two different fields. The qualified field is displayed as label when the axis is zoomed to a deeper level, to show full context.</p> <p>For example, you can generate two fields when showing data by the quarter:</p> <p>A simplified field, with the \$simplified tag, showing quarter, like 'Q1' and a qualified field, with the \$qualified tag, showing year and quarter, like '2016-Q1'.</p> <p>When the time axis is zoomed out, the axis shows labels in two levels, for year (2016) and quarter (Q1), using the simplified field. When you zoom in, the axis shows labels for quarter and month, and the qualified field (2016-Q1) is used to provide full year context for the quarter.</p>
\$cyclic	The \$cyclic tag is used for cyclic fields, for example quarter or month, with a dual data representation.

System fields

In addition to the fields extracted from the data source, system fields are also produced by Qlik Sense. These all begin with "\$" and can be displayed like ordinary fields in a visualization, such as a filter pane or a table. System fields are created automatically when you load data, and are primarily used as an aid in app design.

Available system fields

The following system fields are available:

Available system fields

System field	Description
\$Table	Contains all tables that are loaded.
\$Field	Contains all fields in the tables that are loaded.
\$Fields	Contains the number of fields in each table.
\$FieldNo	Contains the position of the fields in the tables.
\$Rows	Contains the number of rows in the tables

None of the system fields can be manipulated in the script.

Using system fields in a visualization

System field data is associated. For example, if you add two filter panes, one with \$Table and one with \$Field, if you select a table, the \$Field filter pane will show the fields in the selected table as possible values.

System fields are not included in field lists in the assets panel. They are included in the expression editor. If you want to use a system field in the assets panel, you need to reference it by typing it manually.

Example: In a dimension in the assets panel

=*\$Field*

Renaming fields

Sometimes it is necessary to rename fields in order to obtain the desired associations. The three main reasons for renaming fields are:

- Two fields are named differently although they denote the same thing:
 - The field *ID* in the *Customers* table
 - The field *CustomerID* in the *Orders* table

The two fields denote a specific customer identification code and should both be named the same, for example *CustomerID*.

- Two fields are named the same but actually denote different things:
 - The field *Date* in the *Invoices* table
 - The field *Date* in the *Orders* table

The two fields should preferably be renamed, to for example *InvoiceDate* and *OrderDate*.

- There may be errors such as misspellings in the database or different conventions on upper- and lowercase letters.

Since fields can be renamed in the script, there is no need to change the original data. There are two different ways to rename fields as shown in the examples.

Example 1: Using the alias statement

The **LOAD** or **SELECT** statement can be preceded by an **alias** statement.

```
Alias ID as CustomerID;  
LOAD * from Customer.csv;
```

Example 2: Using the as specifier

The **LOAD** or **SELECT** statement can contain the **as** specifier.

```
LOAD ID as CustomerID, Name, Address, Zip, City, State from Customer.csv;
```

Logical tables

Each **LOAD** or **SELECT** statement generates a table. Normally, Qlik Sense treats the result of each one of these as one logical table. However, there are a couple of exceptions from this rule:

- If two or more statements result in tables with identical field names, the tables are concatenated and treated as one logical table.
- If a **LOAD** or **SELECT** statement is preceded by any of the following qualifiers, data is altered or treated differently.

Logical tables

Table	Description
concatenate	This table is concatenated with (added to) another named table or with the last previously created logical table.
crosstable	This table is unpivoted. That is, it is converted from crosstable format to column format.
generic	This table is split into several other logical tables.
info	This table is not loaded as a logical table, but as an information table containing links to external information such as files, sounds, URLs, etc.
intervalmatch	The table (which must contain exactly two columns) is interpreted as numeric intervals, which are associated with discrete numbers in a specified field.
join	This table is joined by Qlik Sense with another named table or with the last previously created logical table, over the fields in common.
keep	This table is reduced to the fields in common with another named table or with the last previously created logical table.
mapping	This table (which must contain exactly two columns) is read as a mapping table, which is never associated with other tables.
semantic	This table is not loaded as a logical table, but as a semantic table containing relationships that should not be joined, e.g. predecessor, successor and other references to other objects of the same type.

When the data has been loaded, the logical tables are associated.

Table names

Qlik Sense tables are named when they are stored in the Qlik Sense database. The table names can be used, for example, for **LOAD** statements with a **resident** clause or with expressions containing the **peek** function, and can be seen in the `$Table` system field in the layout.

Tables are named in accordance with the following rules:

1. If a label immediately precedes a **LOAD** or **SELECT** statement the label is used as table name. The label must be followed by a colon.

Example:

```
Table1:  
LOAD a,b from c.csv;
```

2. If no label is given, the file name or table name immediately following the keyword **FROM** in the **LOAD** or **SELECT** statement is used. A maximum of 32 characters is used. The extension is skipped if the file name is used.
3. Tables loaded inline are named *INLINExx*, where *xx* is a number. The first inline table will be given the name *INLINE01*.
4. Automatically generated tables are named *AUTOGENERATExx*, where *xx* is a number. The first autogenerated table is given the name *AUTOGENERATE01*.
5. If a table name generated according to the rules above should be in conflict with a previous table name, the name is extended with *-x*, where *x* is a number. The number is increased until no conflict remains. For example, three tables could be named *Budget*, *Budget-1* and *Budget-2*.

There are three separate domains for table names: **section access**, **section application** and mapping tables. Table names generated in **section access** and **section application** are treated separately. If a table name referenced is not found within the section, Qlik Sense searches the other section as well. Mapping tables are treated separately and have no connection whatsoever to the other two domains of table names.

Table labels

A table can be labeled for later reference, for example by a **LOAD** statement with a **resident** clause or with expressions containing the **peek** function. The label, which can be an arbitrary string of numbers or characters, should precede the first **LOAD** or **SELECT** statement that creates the table. The label must be followed by a colon ":".

Labels containing blanks must be quoted using single or double quotation marks or square brackets.

Example 1:

```
Table1:  
LOAD a,b from c.csv;  
LOAD x,y from d.csv where x=peek('a',y,'Table1');
```

Example 2: Table label containing a blank

```
[All Transactions]:  
SELECT * from Transtable;  
LOAD Month, sum(Sales) resident [All Transactions] group by Month;
```

Associations between logical tables

A database can have many tables. Each table can be considered as a list of something; each record in the list represents an instance of an object of some type.


Example:

If two tables are lists of different things, for example if one is a list of customers and the other a list of invoices, and the two tables have a field such as the customer number in common, this is usually a sign that there is a relationship between the two tables. In standard SQL query tools the two tables should almost always be joined.

The tables defined in the Qlik Sense script are called logical tables. Qlik Sense makes associations between the tables based on the field names, and performs the joins when a selection is made, for example selecting a field value in a filter pane.

This means that an association is almost the same thing as a join. The only difference is that the join is performed when the script is executed - the logical table is usually the result of the join. The association is made after the logical table is created - associations are always made between the logical tables.

Four tables: a list of countries, a list of customers, a list of transactions and a list of memberships, which are associated with each other through the fields Country and CustomerID.



Qlik Sense association compared to SQL natural outer join

A Qlik Sense association resembles a SQL natural outer join. The association is however more general: an outer join in SQL is usually a one-way projection of one table on another. An association always results in a full (bidirectional) natural outer join.

Frequency information in associating fields

There are some limitations in the use of most associating fields, that is, fields that are common between two or more tables. When a field occurs in more than one table, Qlik Sense has a problem knowing which of the tables it should use for calculating data frequencies.

Qlik Sense analyzes the data to see if there is a non-ambiguous way to identify a main table to count in (sometimes there is), but in most cases the program can only make a guess. Since an incorrect guess could be fatal (Qlik Sense would appear to make a calculation error) the program has been designed not to allow certain operations when the data interpretation is ambiguous for associating fields.

Limitations for associating fields

1. It is not possible to display frequency information in a filter pane showing the field.
2. Statistics boxes for the field show n/a for most statistical entities.
3. In charts, it is not possible to create expressions containing functions that depend on frequency information (such as Sum, Count functions, and Average) on the field, unless the **Distinct** modifier is activated. After each reload, Qlik Sense will scan all chart expressions to see if any ambiguities have occurred as a result of changes in data structures. If ambiguous expressions are found, a warning dialog will be shown and the expression will be disabled. It will not be possible to enable the expression until the problem has been corrected. If a log file is enabled, all ambiguous expressions will be listed in the log.

Workaround

There is a simple way to overcome these limitations. Load the field an extra time under a new name from the table where frequency counts should be made. Then use the new field for a filter pane with frequency, for a statistics box or for calculations in the charts.

Synthetic keys

When two or more data tables have two or more fields in common, this suggests a composite key relationship. Qlik Sense handles this by creating synthetic keys automatically. These keys are anonymous fields that represent all occurring combinations of the composite key.


If you receive a warning about synthetic keys when loading data, it is recommended that you review the data structure in the data model viewer. You should ask yourself whether the data model is correct or not. Sometimes it is, but often enough the synthetic key is there due to an error in the script.

Multiple synthetic keys are often a symptom of an incorrect data model, but not necessarily. However, a sure sign of an incorrect data model is if you have synthetic keys based on other synthetic keys.



When the number of synthetic keys increases, depending on data amounts, table structure and other factors, Qlik Sense may or may not handle them gracefully, and may end up using excessive amount of time and/or memory. In such a case you need to re-work your script by removing all synthetic keys.

Three tables associated with synthetic key \$Syn 1.



Handling synthetic keys

If you need to avoid synthetic keys, there are a number of ways for solving this in the data load script:

- Check that only fields that logically link two tables are used as keys.
 - Fields like “Comment”, “Remark” and “Description” may exist in several tables without being related, and should therefore not be used as keys.
 - Fields like “Date”, “Company” and “Name” may exist in several tables and have identical values, but still have different roles (Order Date/Shipping Date, Customer Company/Supplier Company). In such cases they should not be used as keys.
- Make sure that redundant fields aren't used – that only the necessary fields connect. If for example a date is used as a key, make sure not to load year, month or day_of_month of the same date from more than one internal table.
- If necessary, form your own non-composite keys, typically using string concatenation inside an AutoNumber script function.


Understanding and solving circular references

If there are circular references ("loops") in a data structure, the tables are associated in such a way that there is more than one path of associations between two fields.

This type of data structure should be avoided as much as possible, since it might lead to ambiguities in the interpretation of data.

Three tables with a circular reference, since there is more than one path of associations between two fields.

4 Loading data with the data load script



Qlik Sense solves the problem of circular references by breaking the loop with a loosely coupled table. When Qlik Sense finds circular data structures while executing the load script, a warning dialog will be shown and one or more tables will be set as loosely coupled. Qlik Sense will typically attempt to loosen the longest table in the loop, as this is often a transaction table, which normally should be the one to loosen. In the data model viewer, loosely-coupled tables are indicated by the red dotted links to other tables.

Example:

Data is loaded from three tables that include the following:

- Names of some national soccer teams
- Soccer clubs in some cities
- Cities of some European countries

View of the source data tables in Excel.

The image shows three Excel spreadsheets side-by-side, representing source data tables:

NationalTeams	
Country	Team
Germany	Die Mannschaft
Italy	Azzurri
Spain	La Roja

Clubs	
City	Team
Barcelona	Barcelona
Hamburg	Altona
Madrid	Real Madrid
Milano	Milan
Munich	Bayern München
Rome	Lazio
Turin	Juventus

Cities	
Country	City
Germany	Hamburg
Germany	Munich
Italy	Milano
Italy	Rome
Italy	Turin
Spain	Barcelona
Spain	Madrid

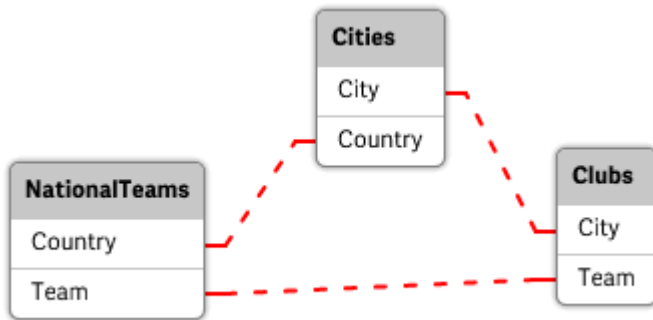
This data structure is not very good, since the field name *Team* is used for two different purposes: national teams and local clubs. The data in the tables creates an impossible logical situation.

When loading the tables into Qlik Sense, Qlik Sense determines which of the data connections that is least important, and loosens this table.

4 Loading data with the data load script

Open the **Data model viewer** to see how Qlik Sense interprets the relevance of the data connections:

View of the circular references as noted by red dotted lines.




The table with cities and the countries they belong to is now loosely coupled to the table with national teams of different countries and to the table with local clubs of different cities.

Solving circular references

When circular references occur, you need to edit the data load script by assigning a unique name to one of the fields with identical names.


Do the following:

1. Open the data load editor.
2. Edit the **LOAD** statement for one of the duplicate field names.
In this example, the **LOAD** statement of the table that holds the local teams and their cities would include with a new name for *Team*, for example *LocalClub*. The updated **LOAD** statement reads:
`LOAD City, Team as LocalClub`
3. Click  in the toolbar to reload data in the app.

You now have logic that works throughout all the tables. In this example, when *Germany* is selected, the national team, the German cities and the local clubs of each city are associated:

Country	Team	City	LocalClub
Germany ✓	Die	Hamburg	Altona
Italy	Azzurri	Munich	Barcelona
Spain	La Roja	Barcelona	
		Madrid	

When you open the **Data model viewer**, you see that the loosely coupled connections are replaced with regular connections:



Concatenating tables

Concatenation is an operation that combines two tables into one.

The two tables are merely added to each other. That is, data is not changed and the resulting table contains the same number of records as the two original tables together. Several concatenate operations can be performed sequentially, so that the resulting table is concatenated from more than two tables.

Automatic concatenation

If the field names and the number of fields of two or more loaded tables are exactly the same, Qlik Sense will automatically concatenate the content of the different statements into one table.

Example:

```
LOAD a, b, c from table1.csv;  
LOAD a, c, b from table2.csv;
```

The resulting internal table has the fields a, b and c. The number of records is the sum of the numbers of records in table 1 and table 2.



The number and names of the fields must be exactly the same. The order of the two statements is arbitrary.

Forced concatenation

Even if two or more tables do not have exactly the same set of fields, it is still possible to force Qlik Sense to concatenate the two tables. This is done with the **concatenate** prefix in the script, which concatenates a table with another named table or with the last previously created table.

Example:

```
LOAD a, b, c from table1.csv;  
concatenate LOAD a, c from table2.csv;
```

The resulting internal table has the fields a, b and c. The number of records in the resulting table is the sum of the numbers of records in table 1 and table 2. The value of field b in the records coming from table 2 is NULL.



Unless a table name of a previously loaded table is specified in the **concatenate** statement the **concatenate** prefix uses the most recently created table. The order of the two statements is thus not arbitrary.

Preventing concatenation

If the field names and the number of fields in two or more loaded tables are exactly the same, Qlik Sense will automatically concatenate the content of the different statements into one table. This is possible to prevent with a **noconcatenate** statement. The table loaded with the associated **LOAD** or **SELECT** statement will then not be concatenated with the existing table.

Example:

```
LOAD a, b, c from table1.csv;  
noconcatenate LOAD a, b, c from table2.csv;
```

Loading data from a previously loaded table

There are two ways to load and transform data from a table that already has been loaded.

- **Resident LOAD** - where you use the **Resident** predicate in a subsequent **LOAD** statement to load a new table.
- Preceding load - where you load from the preceding **LOAD** or **SELECT** statement without specifying a source.

Resident or preceding LOAD?

In most cases, the same result can be achieved by using either method. A preceding **LOAD** is generally the faster option, but there are some cases where you need to use a **Resident LOAD** instead:

- If you want to use the **Order_by** clause to sort the records before processing the **LOAD** statement.
- If you want to use any of the following prefixes, in which cases preceding **LOAD** is not supported:
 - **Crosstable**
 - **Join**
 - **Intervalmatch**

Resident LOAD

You can use the **Resident** predicate in a **LOAD** statement to load data from a previously loaded table. This is useful when you want to perform calculations on data loaded with a **SELECT** statement where you do not have the option to use Qlik Sense functions, such as date or numeric value handling.

Example:

In this example, the date interpretation is performed in the **Resident** load as it can't be done in the initial **Crosstable LOAD**.

```
PreBudget:  
Crosstable (Month, Amount, 1)
```

```
LOAD Account,
    Jan,
    Feb,
    Mar,
...
From Budget;

Budget:
Noconcatenate
LOAD
    Account,
    Month(Date#(Month, 'MMM')) as Month,
    Amount
Resident PreBudget;

Drop Table PreBudget;
```



A common case for using **Resident** is where you want to use a temporary table for calculations or filtering. Once you have achieved the purpose of the temporary table, it should be dropped using the **Drop table** statement.

Preceding load

The preceding load feature allows you to load a table in one pass, but still define several successive transformations. Basically, it is a **LOAD** statement that loads from the **LOAD** or **SELECT** statement below, without specifying a source qualifier such as **From** or **Resident** as you would normally do. You can stack any number of **LOAD** statements this way. The statement at the bottom will be evaluated first, then the statement above, and so on until the top statement has been evaluated.

You can achieve the same result using **Resident**, but in most cases a preceding **LOAD** will be faster.

Another advantage of preceding load is that you can keep a calculation in one place, and reuse it in **LOAD** statements placed above.



The following prefixes cannot be used in conjunction with preceding **LOAD:Join, Crosstable** and **Intervalmatch**.

Example 1: Transforming data loaded by a **SELECT** statement

If you load data from a database using a **SELECT** statement, you cannot use Qlik Sense functions to interpret data in the **SELECT** statement. The solution is to add a **LOAD** statement, where you perform data transformation, above the **SELECT** statement.

In this example we interpret a date stored as a string using the Qlik Sense function **Date#** in a **LOAD** statement, using the previous **SELECT** statement as source.

```
LOAD Date#(OrderDate, 'YYYYMMDD') as orderDate;
SQL SELECT OrderDate FROM ... ;
```

Example 2: Simplifying your script by reusing calculations

In this example we use a calculation more than once in the script:

```
LOAD ...,
    Age( FromDate + IterNo() - 1, BirthDate ) as Age,
    Date( FromDate + IterNo() - 1 ) as ReferenceDate
Resident Policies
    while IterNo() <= ToDate - FromDate + 1 ;
```

By introducing the calculation in a first pass, we can reuse it in the Age function in a preceding **LOAD**:

```
LOAD ..., ReferenceDate,
    Age( ReferenceDate, BirthDate ) as Age;
LOAD *,
    Date( FromDate + IterNo() - 1 ) as ReferenceDate
Resident Policies
    while IterNo() <= ToDate - FromDate + 1 ;
```

Data types in Qlik Sense

Qlik Sense can handle text strings, numbers, dates, times, timestamps, and currencies correctly. They can be sorted, displayed in a number of different formats and they can be used in calculations. This means, for example, that dates, times, and timestamps can be added to or subtracted from each other.

Data representation inside Qlik Sense

In order to understand data interpretation and number formatting in Qlik Sense, it is necessary to know how data is stored internally by the program. All of the data loaded into Qlik Sense is available in two representations: as a string and as a number.

1. The string representation is always available and is what is shown in the list boxes and the other sheet objects. Formatting of data in list boxes (number format) only affects the string representation.
2. The number representation is only available when the data can be interpreted as a valid number. The number representation is used for all numeric calculations and for numeric sorting.

If several data items read into one field have the same number representation, they will all be treated as the same value and will all share the first string representation encountered. Example: The numbers 1.0, 1 and 1.000 read in that order will all have the number representation 1 and the initial string representation 1.0.

Number interpretation

When you load data containing numbers, currency, or dates, it will be interpreted differently depending on whether the data type is defined or not. This section describes how data is interpreted in the two different cases.

Data with type information

Fields containing numbers with a defined data type in a database loaded using ODBC will be handled by Qlik Sense according to their respective formats. Their string representation will be the number with an appropriate formatting applied.

Qlik Sense will remember the original number format of the field even if the number format is changed for a measure under **Number formatting** in the properties panel.

4 Loading data with the data load script

The default formats for the different data types are:

- integer, floating point numbers: the default setting for number
- currency: the default setting for currency
- time, date, timestamp: ISO standard formatting

The default settings for number and currency are defined using the script number interpretation variables or the operating system settings (**Control Panel**).

Data without type information

For data without specific formatting information from the source (for example, data from text files or ODBC data with a general format) the situation becomes more complicated. The final result will depend on at least six different factors:

- The way data is written in the source database
- The operating system settings for number, time, date and so on. (**Control Panel**)
- The use of optional number-interpreting variables in the script
- The use of optional interpretation functions in the script
- The use of optional formatting functions in the script
- The number formatting controls in the document

Qlik Sense tries to interpret input data as a number, date, time, and so on. As long as the system default settings are used in the data, the interpretation and the display formatting is done automatically by Qlik Sense, and the user does not need to alter the script or any setting in Qlik Sense.

By default, the following scheme is used until a complete match is found. (The default format is the format such as the decimal separator, the order between year, month and day, and so on, specified in the operating system, that is, in the **Control Panel**, or in some cases from the special number interpretation variables in the script.

Qlik Sense will interpret the data as:

1. A number in accordance with the default format for numbers.
2. A date according to the default format for date.
3. A timestamp according to the default format for time and date.
4. A time according to the default format for time.
5. A date according to the following format: yyyy-MM-dd.
6. A time-stamp according to the following format: YYYY-MM-DD hh:mm[:ss[.fff]].
7. A time according to the following format: hh:mm[:ss[.fff]].
8. Money according to the default format for currency.
9. A number with '.' as decimal separator and ',' as thousands separator, provided that neither the decimal separator nor the thousands separator are set to ','.
10. A number with ',' as decimal separator and '.' as thousands separator, provided that neither the decimal separator nor the thousands separator are set to '.'.
11. A text string. This last test never fails: if it is possible to read the data, it is always possible to interpret it as a string.

When loading numbers from text files, some interpretation problems may occur, for example, an incorrect thousands separator or decimal separator may cause Qlik Sense to interpret the number incorrectly. The first thing to do is to check that the number-interpretation variables in the script are correctly defined and that the system settings in the **Control Panel** are correct.

When Qlik Sense has interpreted data as a date or time, it is possible to change to another date or time format in the properties panel of the visualization.

Since there is no predefined format for the data, different records may, of course, contain differently formatted data in the same field. It is possible for example, to find valid dates, integers, and text in one field. The data will therefore, not be formatted, but shown in its original form.

Date and time interpretation

Qlik Sense stores each date, time, and timestamp found in data as a date serial number. The date serial number is used for dates, times and timestamps and in arithmetic calculations based on date and time entities. Dates and times can thus be added and subtracted, intervals can be compared, and so on.

The date serial number is the (real valued) number of days passed since December 30, 1899, that is, the Qlik Sense format is identical to the 1900 date system used by Microsoft Excel and other programs, in the range between March 1, 1900 and February 28, 2100. For example, 33857 corresponds to September 10, 1992. Outside this range, Qlik Sense uses the same date system extended to the Gregorian calendar.

The serial number for times is a number between 0 and 1. The serial number 0.00000 corresponds to 00:00:00, whereas 0.99999 corresponds to 23:59:59. Mixed numbers indicate the date and time: the serial number 2.5 represents January 1, 1900 at 12:00 noon.

The data is, however, displayed according to the format of the string. By default, the settings made in the **Control Panel** are used. It is also possible to set the format of the data by using the number interpretation variables in the script or with the help of a formatting function. Lastly, it is also possible to reformat the data in the properties sheet of the sheet object.

Example 1:

- 1997-08-06 is stored as 35648
- 09:00 is stored as 0.375
- 1997-08-06 09:00 is stored as 35648.375

and the other way around:

- 35648 with number format 'D/M/YY' is shown as 6/8/97
- 0.375 with number format 'hh.mm' is shown as 09.00

Qlik Sense follows a set of rules to try to interpret dates, times, and other data types. The final result, however, will be affected by a number of factors as described here.

Example 2:

These examples assume the following default settings:

4 Loading data with the data load script

- Number decimal separator:
- Short date format: YY-MM-DD
- Time format: hh:mm

The following table shows the different representations when data is read into Qlik Sense without the special interpretation function in the script:

Table when data is read without the special interpretation function in the script

Source data	Qlik Sense default interpretation	'YYYY-MM-DD' date format	'MM/DD/YYYY' date format	'hh:mm' time format	'# ##0.00' number format
0.375	0.375	1899-12-30	12/30/1899	09:00	0.38
33857	33857	1992-09-10	09/10/1992	00:00	33 857.00
97-08-06	97-08-06	1997-08-06	08/06/1997	00:00	35 648.00
970806	970806	4557-12-21	12/21/4557	00:00	970 806.00
8/6/97	8/6/97	8/6/97	8/6/97	8/6/97	8/6/97

The following table shows the different representations when data is read into Qlik Sense using the date#(A, 'M/D/YY') interpretation function in the script:

Table when using the date#(A, 'M/D/YY') interpretation function in the script

Source data	Qlik Sense default interpretation	'YYYY-MM-DD' date format	'MM/DD/YYYY' date format	'hh:mm' time format	'# ##0.00' number format
0.375	0.375	0.375	0.375	0.375	0.375
33857	33857	33857	33857	33857	33857
97-08-06	97-08-06	97-08-06	97-08-06	97-08-06	97-08-06
970806	970806	970806	970806	970806	970806
8/6/97	8/6/97	1997-08-06	08/06/1997	00:00	35 648.00

Dollar-sign expansions

Dollar-sign expansions are definitions of text replacements used in the script or in expressions. This process is known as expansion - even if the new text is shorter. The replacement is made just before the script statement or the expression is evaluated. Technically it is a macro expansion.

The expansion always begins with '\$(' and ends with ') ' and the content between brackets defines how the text replacement will be done. To avoid confusion with script macros we will henceforth refer to macro expansions as dollar-sign expansions.

Dollar-sign expansions can be used with either of:

- variables
- parameters
- expressions



A dollar-sign expansion is limited in how many expansions it can calculate. Any expansion with over 1000 levels of nested expansions will not be calculated.

Dollar-sign expansion using a variable

When using a variable for text replacement in the script or in an expression, the following syntax is used:

```
$(variablename)
```

`$(variablename)` expands to the value in the variable. If `variablename` does not exist, the expansion will result in an empty string.

For numeric variable expansions, the following syntax is used:

```
$(#variablename)
```

It always yields a valid decimal-point representation of the numeric value of the variable, possibly with exponential notation (for very large/small numbers). If `variablename` does not exist or does not contain a numeric value, it will be expanded to `0` instead.

Example:

```
SET DecimalSep=',';  
LET X = 7/2;
```

The dollar-sign expansion `$(X)` will expand to `3,5` while `$(#X)` will expand to `3.5`.

Example:

```
Set MyPath=C:\MyDocs\Files\  
...  
LOAD * from $(MyPath)abc.csv;  
Data will be loaded from C:\MyDocs\Files\abc.csv.
```

Example:

```
Set CurrentYear=1992;  
...  
SQL SELECT * FROM table1 WHERE Year=$(CurrentYear);  
Rows with Year=1992 will be selected.
```

Example:

```
Set vConcatenate = ;  
For each vFile in FileList('.\*.txt')  
  Data:  
    $(vConcatenate)  
    LOAD * FROM [$(vFile)];
```

4 Loading data with the data load script

```
Set vConcatenate = Concatenate ;
Next vFile
```

In this example, all .txt files in the directory are loaded using the **Concatenate** prefix. This may be required if the fields differ slightly, in which case auto-concatenation does not work. The vConcatenate variable is initially set to an empty string, as the **Concatenate** prefix cannot be used on the first load. If the directory contains three files named *file1.txt*, *file2.txt* and *file3.txt*, the **LOAD** statement will during the three iterations expand to:

```
LOAD * FROM[.\file1.txt];
Concatenate LOAD * FROM[.\file2.txt];
Concatenate LOAD * FROM[.\file3.txt];
```

Dollar-sign expansion using parameters

Parameters can be used in dollar-sign expansions. The variable must then contain formal parameters, such as \$1, \$2, \$3 etc. When expanding the variable, the parameters should be stated in a comma separated list.

Example:

```
Set MUL='$1*$2';
Set X=$(MUL(3,7)); // returns '3*7' in X

Let X=$(MUL(3,7)); // returns 21 in X
```

If the number of formal parameters exceeds the number of actual parameters only the formal parameters corresponding to actual parameters will be expanded. If the number of actual parameters exceeds the number of formal parameters the superfluous actual parameters will be ignored.

Example:

```
Set MUL='$1*$2';
Set X=$(MUL); // returns '$1*$2' in X

Set X=$(MUL(10)); // returns '10*$2' in X

Let X=$(MUL(5,7,8)); // returns 35 in X
```

The parameter \$0 returns the number of parameters actually passed by a call.

Example:

```
set MUL='$1*$2 $0 par';
set X=$(MUL(3,7)); // returns '3*7 2 par' in X
```

Dollar-sign expansion using an expression

Expressions can be used in dollar-sign expansions. The content between the brackets must then start with an equal sign:

```
$(=expression )
```

The expression will be evaluated and the value will be used in the expansion.

Example:

```
$(=Year(Today())); // returns a string with the current year.
```

```
$(=only(Year)-1); // returns the year before the selected one.
```

File inclusion

File inclusions are made using dollar-sign expansions. The syntax is:

```
$(include=filename )
```

The above text will be replaced by the content of the file specified after the equal sign. This feature is very useful when storing scripts or parts of scripts in text files.

Example:

```
$(include=C:\Documents\MyScript.qvs);
```

Using quotation marks in the script

You can use quotation marks in script statements in a number of different ways.

Inside LOAD statements

In a **LOAD** statement the following symbols should be used as quotation marks:

Field names

Description	Symbol	Code point	Example
double quotation marks	" "	34	"string"
square brackets	[]	91, 93	[string]
grave accents	` `	96	`string`

String literals

Description	Symbol	Code point	Example
single quotation marks	' '	39	'string'

In SELECT statements

For a **SELECT** statement interpreted by an ODBC driver, usage may vary. Usually, you should use the straight double quotation marks (Alt + 0034) for field and table names, and the straight single quotation marks (Alt + 0039) for literals, and avoid using grave accents. However, some ODBC drivers not only accept grave accents as quotation marks, but also prefer them. In such a case, the generated **SELECT** statements contain grave accent quotation marks.

Microsoft Access quotation marks example

Microsoft Access ODBC Driver 3.4 (included in Microsoft Access 7.0) accepts the following quotation marks when analyzing the **SELECT** statement:

Field names and table names:

- []
- " "
- ` `

String literals:

- ''

Other databases may have different conventions.

Outside LOAD statements

Outside a **LOAD** statement, in places where Qlik Sense expects an expression, double quotation marks denote a variable reference and not a field reference. If you use double quotation marks, the enclosed string will be interpreted as a variable and the value of the variable will be used.

Out-of-context field references and table references

Some script functions refer to fields that have already been created, or are in the output of a **LOAD** statement, for example **Exists()** and **Peek()**. These field references are called out-of-context field references, as opposed to source field references that refer to fields that are in context, that is, in the input table of the **LOAD** statement.

Out-of-context field references and table references should be regarded as literals and therefore need single quotation marks.

Difference between names and literals

The difference between names and literals becomes clearer comparing the following examples:

Example:

'Sweden' as Country

When this expression is used as a part of the field list in a **LOAD** or **SELECT** statement, the text string "Sweden" will be loaded as field value into the Qlik Sense field "Country".

Example:

"land" as Country

When this expression is used as a part of the field list in a **LOAD** or **SELECT** statement, the content of the database field or table column named "land" will be loaded as field values into the Qlik Sense field "Country". This means that *land* will be treated as a field reference.

Difference between numbers and string literals

The difference between numbers and string literals becomes clearer comparing the following examples.

Example:

```
'12/31/96'
```

When this string is used as a part of an expression, it will in a first step be interpreted as the text string "12/31/96", which in turn may be interpreted as a date if the date format is 'MM/DD/YY'. In that case it will be stored as a dual value with both a numeric and a textual representation.

Example:

```
12/31/96
```

When this string is used as a part of an expression, it will be interpreted numerically as 12 divided by 31 divided by 96.

Using quotation marks in a string

When a string contains characters that can be used as quotation marks, it is important to clearly indicate where a string begins and where it ends when quoting the string. If the string is not properly quoted, the script will fail or it will load data incorrectly.

There are two methods for quoting a string that contains quotation marks.

Use a specific quotation mark to quote the string

Choose a quotation mark that is not used inside the string, and use it to quote the entire string. Qlik Sense will use that specific quotation mark to determine where the string begins and ends.

Any of the following quotation marks can be used to quote the entire string:

- Double quotation marks " "
- Square brackets []
- Grave accents ` `
- Single quotation marks ' '

Example:

```
[Table '1 "2"]
```

Square brackets are used to quote the string. The string loads as: *Table '1 "2"*

```
'string `Name1` "Name2"
```

Single quotation marks are used to quote the string. The string loads as: *string `Name1` "Name2"*

Use escape characters

Escape characters are an additional instance of the quotation mark that is used to quote the string. They must be added beside every instance of the quotation mark that appears within the string. When all quotation marks

4 Loading data with the data load script

are used inside a string, you need to add escape characters beside the same type of quotation mark used to quote the string. Escape characters can also be used if you want to use a quotation mark that is already in use in a string.

Only the following marks can be used as escape characters:

- Double quotation marks " "
- Square brackets []
- Single quotation marks ' '

Example:

```
"Michael said ""It's a beautiful day""."
```

If you quote the string using the double quotation marks " ", then you must add an extra double quotation mark beside every double quotation mark used inside the string.

This string is loaded as *Michael said "It's a beautiful day"*. By using the escape character "", the Qlik Sense Data load editor understands which double quotation marks are part of the string and which quotation mark indicates the end of the string. The single quotation mark ' used in the abbreviation *It's* does not need to be escaped because it is not the mark used to quote the string.

Example:

```
'Michael said: "It''s a beautiful day".'
```

If you quote this string using single quotation marks, then you must add an extra single quotation mark beside each single quotation mark used inside the string.

This string is loaded as *Michael said "It's a beautiful day"*. The double quotation mark " used for quoting what Michael said does not need to be escaped because it is not the mark used to quote the string.

Example:

```
[Michael said [It's a "beautiful day]].]
```

Square brackets [] behave differently from the other two quotation marks. If you want to use brackets as an escape character, you must add an extra bracket beside the right square bracket] only, and not beside the left square bracket [.

This string is loaded as *Michael said [It's a "beautiful day]*. Only the right square bracket] is escaped. The single quotation mark ' and the double quotation mark " used in the string do not need to be escaped as they are not used to quote the string.

Wild cards in the data

It is also possible to use wild cards in the data. Two different wild cards exist: the star symbol, interpreted as all values of this field, and an optional symbol, interpreted as all remaining values of this field.

The star symbol

The star symbol is interpreted as all (listed) values of this field, that is, a value listed elsewhere in this table. If used in one of the system fields (*USERID*, *PASSWORD*, *NTNAME* or *SERIAL*) in a table loaded in the access section of the script, it is interpreted as all (also not listed) possible values of this field.

The star symbol is not allowed in information files. Also, it cannot be used in key fields (that is, fields used to join tables).

There is no star symbol available unless explicitly specified. .

OtherSymbol

In many cases a way to represent all other values in a table is needed, that is, all values that were not explicitly found in the loaded data. This is done with a special variable called **OtherSymbol**. To define the **OtherSymbol** to be treated as "all other values", use the following syntax:

```
SET OTHERSYMBOL=<sym>;
```

before a **LOAD** or **SELECT** statement. <sym> may be any string.

The appearance of the defined symbol in an internal table will cause Qlik Sense to define it as all values not previously loaded in the field where it is found. Values found in the field after the appearance of the **OtherSymbol** will thus be disregarded.

In order to reset this functionality use:

```
SET OTHERSYMBOL=;
```

Example:

Table Customers

CustomerID	Name
1	ABC Inc.
2	XYZ Inc.
3	ACME INC
+	Undefined

Table Orders

CustomerID	Name
1	1234
3	1243
5	1248
7	1299

Insert the following statement in the script before the point where the first table above is loaded:

4 Loading data with the data load script

```
SET OTHERSYMBOL=+;
```

Any reference to a *CustomerID* other than 1, 2 or 3, e.g. as when clicking on *OrderID 1299* will result in *Undefined* under *Name*.



OtherSymbol is not intended to be used for creating outer joins between tables.

NULL value handling

When no data can be produced for a certain field as a result of a database query and/or a join between tables, the result is normally a NULL value.

Overview

The Qlik Sense logic treats the following as real NULL values:

- NULL values returned from an ODBC connection
- NULL values created as a result of a forced concatenation of tables in the data load script
- NULL values created as a result of a join made in the data load script
- NULL values created as a result of the generation of field value combinations to be displayed in a table



It is generally impossible to use these NULL values for associations and selections, except when the **NullAsValue** statement is being employed.

Text files per definition cannot contain NULL values.

Associating/selecting NULL values from ODBC

It is possible to associate and/or select NULL values from an ODBC data source. For this purpose a script variable has been defined. The following syntax can be used:

```
SET NULLDISPLAY=<sym>;
```

The symbol <sym> will substitute all NULL values from the ODBC data source on the lowest level of data input. <sym> may be any string.

In order to reset this functionality to the default interpretation, use the following syntax:

```
SET NULLDISPLAY=;
```



The use of **NULLDISPLAY** only affects data from an ODBC data source.

If you wish to have the Qlik Sense logic interpret NULL values returned from an ODBC connection as an empty string, add the following to your script before any **SELECT** statement:

```
SET NULLDISPLAY="";
```



Here " is actually two single quotation marks without anything in between.

Creating NULL values from text files

It is possible to define a symbol, which when it occurs in a text file or an **inline** clause will be interpreted as a real NULL value. Use the following statement:

```
SET NULLINTERPRET=<sym>;
```

The symbol <sym> is to be interpreted as NULL. <sym> may be any string.

In order to reset this functionality to the default interpretation, use:

```
SET NULLINTERPRET=;
```



*The use of **NULLINTERPRET** only affects data from text files and inline clauses.*

Propagation of NULL values in expressions

NULL values will propagate through an expression according to a few logical and quite reasonable rules.

Functions

The general rule is that functions return NULL when the parameters fall outside the range for which the function is defined.

Examples

Expression	Result
asin(2)	returns NULL
log(-5)	returns NULL
round(A,0)	returns NULL

As a result of the above follows that functions generally return NULL when any of the parameters necessary for the evaluation are NULL.

Examples

Expression	Result
sin(NULL)	returns NULL
chr(NULL)	returns NULL
if(NULL, A, B)	returns B
if(True, NULL, A)	returns NULL
if(True, A, NULL)	returns A

The exception to the second rule are logical functions testing for type.

4 Loading data with the data load script

Examples

Expression	Result
isnull(NULL)	returns True (-1)
isnum(NULL)	returns False (0)

Arithmetic and string operators

If NULL is encountered on any side of these operators NULL is returned.

Examples

Expression	Result
A + NULL	returns NULL
A - NULL	returns NULL
A / NULL	returns NULL
A * NULL	returns NULL
NULL / A	returns NULL
0 / NULL	returns NULL
0 * NULL	returns NULL
A & NULL	returns A

Relational operators

If NULL is encountered on any side of relational operators special rules apply.

Examples

Expression	Result
NULL (any relation operator) NULL	returns NULL
A <> NULL	returns True (-1)
A < NULL	returns False (0)
A <= NULL	returns False (0)
A = NULL	returns False (0)
A >= NULL	returns False (0)
A > NULL	returns False (0)

4.4 Guidelines for data and fields

There are certain conventions and limitations you need to be aware of when working with Qlik Sense. For example: the upper limit for data tables and fields as well as maximum amount of loaded data in Qlik Sense.

Guidelines for amount of loaded data

The amount of data that can be loaded into Qlik Sense is primarily limited by the amount of primary memory of the computer.

Upper limits for data tables and fields

Be aware when building very large apps that a Qlik Sense app cannot have more than 2,147,483,648 distinct values in one field.

The number of fields and data tables as well as the number of table cells and table rows that can be loaded, is limited only by RAM.

Recommended limit for load script sections

The recommended maximum number of characters to be used per load script section is 50,000 characters.

Conventions for number and time formats

In many interpretation and formatting functions it is possible to set the format for numbers and dates by using a format code. This topic describes the conventions used to format a number, date, time, or time stamp. These conventions apply both to script and chart functions.

Number formats

To denote a specific number of digits, use the symbol "0" for each digit.

To denote a possible digit to the left of the decimal point, use the symbol "#".

To mark the position of the thousands separator or the decimal separator, use the applicable thousands separator and the decimal separator.

The format code is used for defining the positions of the separators. It is not possible to set the separator in the format code. Use the **DecimalSep** and **ThousandSep** variables for this in the script.

It is possible to use the thousand separator to group digits by any number of positions, for example, a format string of "0000-0000-0000" (thousand separator="-") could be used to display a twelve-digit part number as "0012-4567-8912".

Examples:

Example of number formats

Number format	Description
# ##0	describes the number as an integer with a thousands separator. In this example " " is used as a thousands separator.
###0	describes the number as an integer without a thousands separator.

4 Loading data with the data load script

Number format	Description
0000	describes the number as an integer with at least four digits. For example, the number 123 will be shown as 0123.
0.000	describes the number with three decimals. In this example "." is used as a decimal separator.

Special number formats

Qlik Sense can interpret and format numbers in any radix between 2 and 36 including binary, octal and hexadecimal. It can also handle roman formats.

Special number formats

Format	Description
Binary format	To indicate binary format the format code should start with (bin) or (BIN).
Octal format	To indicate octal format the format code should start with (oct) or (OCT).
Hexadecimal format	To indicate hexadecimal format the format code should start with (hex) or (HEX). If the capitalized version is used A-F will be used for formatting (for example 14FA). The non-capitalized version will result in formatting with a-f (for example 14fa). Interpretation will work for both variants regardless of the capitalization of the format code.
Decimal format	The use of (dec) or (DEC) to indicate decimal format is permitted but unnecessary.
Custom radix format	To indicate a format in any radix between 2 and 36 the format code should start with (rxx) or (Rxx) where xx is the two-digit number denoting the radix to be used. If the capitalized R is used letters in radices above 10 will be capitalized when Qlik Sense is formatting (for example 14FA). The non-capitalized r will result in formatting with non-capital letters (for example 14fa). Interpretation will work for both variants regardless of the capitalization of the format code. Note that (r02) is the equivalent of (bin), (R16) is the equivalent of (HEX), and so on.
Roman format	To indicate roman numbers the format code should start with (rom) or (ROM). If the capitalized version is used capital letters will be used for formatting (for example MMXVI). The non-capitalized version will result in formatting with lower cap letters (mmxvi). Interpretation will work for both variants regardless of the capitalization of the format code. Roman numbers are generalized with minus sign for negative numbers and 0 for zero. Decimals are ignored with roman formatting.

Examples:

Examples of special number formats

Example	Result
num(199, '(bin)')	returns 11000111
num(199, '(oct)')	returns 307

4 Loading data with the data load script

Example	Result
num(199, '(hex)')	returns c7
num(199, '(HEX)')	returns C7
num(199, '(r02)')	returns 11000111
num(199, '(r16)')	returns c7
num(199, '(R16)')	returns C7
num(199, '(R36)')	returns 5J
num(199, '(rom)')	returns cxcix
num(199, '(ROM)')	returns CXCIX

Dates

You can use the following symbols to format a date. Arbitrary separators can be used.

Symbols to format a date

Symbols	Description
D	To describe the day, use the symbol "D" for each digit.
M	To describe the month number, use the symbol "M". Use "M" or "MM" for one or two digits. "MMM" denotes short month name in letters as defined by the operating system or by the override system variable MonthNames in the script. "MMMM" denotes long month name in letters as defined by the operating system or by the override system variable LongMonthNames in the script.
Y	To describe the year, use the symbol "Y" for each digit.
W	To describe the weekday, use the symbol "W". "W" will return the number of the day (for example 0 for Monday) as a single digit. "WW" will return the number with two digits (e.g. 02 for Wednesday). "WWW" will show the short version of the weekday name (for example Mon) as defined by the operating system or by the override system variable DayNames in the script. "WWWW" will show the long version of the weekday name (for example Monday) as defined by the operating system or by the override system variable LongDayNames in the script.

Examples: (with 31st March 2013 as example date)

Examples of date formats

Example	Result
YY-MM-DD	describes the date as 13-03-31.
YYYY-MM-DD	describes the date as 2013-03-31.
YYYY-MMM-DD	describes the date as 2013-Mar-31.
DD MMMM YYYY	describes the date as 31 March 2013.
M/D/YY	describes the date as 3/31/13.
W YY-MM-DD	describes the date as 6 13-03-31.
WWW YY-MM-DD	describes the date as Sat 13-03-31.
WWWW YY-MM-DD	describes the date as Saturday 13-03-31.

Times

You can use the following symbols to format a time. Arbitrary separators can be used.

Symbols to format a time

Symbols	Description
h	To describe the hours, use the symbol "h" for each digit.
m	To describe the minutes, use the symbol "m" for each digit.
s	To describe the seconds, use the symbol "s" for each digit.
f	To describe the fractions of a second, use the symbol "f" for each digit.
tt	To describe the time in AM/PM format, use the symbol "tt" after the time.

Examples: (with 18.30 as example time):

Examples of time formats

Example	Result
hh:mm	describes the time as 18:30.
hh.mm.ss.ff	describes the time as 18.30.00.00
hh:mm:tt	describes the time as 18.30.00.00.

Time stamps

The same notation as that of dates and times above is used in time stamps.

Examples: (with 31th March 2013 18.30 as example time stamp):

Examples of time stamp formats

Example	Result
YY-MM-DD hh:mm	describes the time stamp as 13-03-31 18:30.
M/D/Y hh.mm.ss.ffff	describes the time stamp as 3/31/13 18.30.00.0000.

4.5 Working with QVD files

A QVD (QlikView Data) file is a file containing a table of data exported from Qlik Sense. QVD is a native Qlik format and can only be written to and read by Qlik Sense or QlikView. The file format is optimized for speed when reading data from a script but it is still very compact. Reading data from a QVD file is typically 10-100 times faster than reading from other data sources.

QVD files can be read in two modes: standard (fast) and optimized (faster). The selected mode is determined automatically by the script engine.

There are some limitations regarding optimized loads. It is possible to rename fields, but any of the operations mentioned here will disable the optimized load and result in a standard load.

- Any transformations on the fields that are loaded.
- Using a **where** clause causing Qlik Sense to unpack the records.
- Using **Map** on a field that is loaded.

Purpose of QVD files

QVD files can be used for many purposes and more than one may apply in any given situation. At least four major uses can be easily identified:

- Increasing load speed
By buffering non-changing or slowly-changing blocks of input data in QVD files, script execution becomes considerably faster for large data sets.
- Decreasing load on database servers
The amount of data fetched from external data sources can also be greatly reduced. This reduces the workload on external databases and network traffic. Furthermore, when several scripts share the same data, it is only necessary to load it once from the source database into a QVD file. Other apps can make use of the same data through this QVD file.
- Consolidating data from multiple apps
With the **binary** script statement, data can be loaded from a single app into another app, but with QVD files a script can combine data from any number of apps. This makes it possible for apps to consolidate similar data from different business units, for example.
- Incremental

In many common cases, the QVD functionality can be used for incremental load by loading only new records from a growing database.

Creating QVD files

A QVD file can be created in two ways:

- Explicit creation and naming using the **store** command in the script. State in the script that a previously-read table, or part thereof, is to be exported to an explicitly-named file at a location of your choice.
- Automatic creation and maintenance from script. When you precede a **LOAD** or **SELECT** statement with the **buffer** prefix, Qlik Sense will automatically create a QVD file, which, under certain conditions, can be used instead of the original data source when reloading data.

There is no difference between the resulting QVD files with regard to reading speed.

Reading data from QVD files

A QVD file can be read or accessed by the following methods:

- Loading a QVD file as an explicit data source. QVD files can be referenced by a **LOAD** statement in the script, just like any other type of text files (csv, fix, dif, biff etc).

For example (Windows):

- `LOAD * from xyz.qvd (qvd)`
- `LOAD Name, RegNo from xyz.qvd (qvd)`
- `LOAD Name as a, RegNo as b from xyz.qvd (qvd)`

For example (Kubernetes):

- `LOAD * from [lib://MyDataFiles/xyz.qvd];`
- `LOAD Name, RegNo from [lib://MyDataFiles/xyz.qvd];`
- `LOAD Name as a, RegNo as b from [lib://MyDataFiles/xyz.qvd];`

- Automatic loading of buffered QVD files. When you use the **buffer** prefix on **LOAD** or **SELECT** statements, no explicit statements for reading are necessary. Qlik Sense will determine the extent to which it will use data from the QVD file as opposed to acquiring data using the original **LOAD** or **SELECT** statement.
- Accessing QVD files from the script. A number of script functions (all beginning with **qvd**) can be used for retrieving various information on the data found in the XML header of a QVD file.

QVD format

A QVD file holds exactly one data table and consists of three parts:

- Header.



If the QVD file was generated with QlikView the header is a well-formed XML header (in UTF-8 char set) describing the fields in the table, the layout of the subsequent information and other metadata.

- Symbol tables in a byte-stuffed format.
- Actual table data in a bit-stuffed format.

4.6 Managing security with section access

You can use section access in the data load script to handle security. In this way, a single file can be used to hold the data for a number of users or user groups. Qlik Sense uses the information in the section access for authentication and authorization, and dynamically reduces the data, so that users only see their own data.



User groups are not supported in Kubernetes.

The security is built into the file itself, which means a downloaded file is also protected, to some extent. However, if security demands are high, downloading of files and offline use should be prevented, and files should be published by the Qlik Sense server only. As all data is kept in one file, the size of this file can potentially be very large.



To avoid exposing restricted data, remove all attached files with section access settings before publishing the app.

Attached files are included when the app is published. If the published app is copied, the attached files are included in the copy. However, if section access restrictions have been applied to the attached data files, the section access settings are not retained when the files are copied, so users of the copied app will be able to see all the data in the attached files.



A snapshot shows data according to the access rights of the user who takes the snapshot, and the snapshot can then be shared in a story. However, when users return to a visualization from a story to see the live data in the app, they are restricted by their own access rights.



You must not assign colors to master dimension values if you use section access or work with sensitive data because the values may be exposed.

Sections in the script

Access control is managed through one or several security tables loaded in the same way that Qlik Sense normally loads data. This makes it possible to store these tables in a normal database. The script statements managing the security tables are given within the access section, which in the script is initiated by the statement **Section Access**.

If an access section is defined in the script, the part of the script loading the app data must be put in a different section, initiated by the statement **Section Application**.

Example:

```
Section Access;  
LOAD * inline  
[ACCESS,USERID  
USER,U ];  
Section Application;  
LOAD... .. from... ..
```

Section access system fields

The access levels are assigned to users in one or several tables loaded within the section access. These tables can contain several different user-specific system fields, typically USERID, and the field defining the access level, ACCESS. All section access system fields will be used for authentication or authorization. The full set of section access system fields is described below.

ACCESS

Defines what access the corresponding user should have.

Access to Qlik Sense apps can be authorized for specified users or groups of users. In the security table, users can be assigned to the access levels ADMIN or USER. If no valid access level is assigned, the user cannot open the app.

A person with ADMIN privileges has access to all data in the app. A person with USER privileges can only access data as defined in the security table.

USERID

Contains a string corresponding to a Qlik Sense user name. Qlik Sense will get the login information from the proxy and compare it to the value in this field.

GROUP

Contains a string corresponding to a group in Qlik Sense. Qlik Sense will resolve the user supplied by the proxy against this group.



When you use groups to reduce data and want to Qlik Management Console, the INTERNAL\SA_SCHEDULER account user is still required.

OMIT

Contains the name of the field that is to be omitted for this specific user. Wildcards may be used and the field may be empty. An easy way of doing this is to use a sub field.



We recommend that you do not apply OMIT on key fields. Key fields that are omitted are visible in the data model viewer, but the content is not available, which can be confusing for a user. Additionally, applying OMIT on fields that are used in a visualization can result in an incomplete visualization for users that do not have access to the omitted fields.

4 Loading data with the data load script

Qlik Sense will compare the user supplied by the proxy with UserID and resolve the user against groups in the table. If the user belongs to a group that is allowed access, or the user matches, they will get access to the app.



If you have locked yourself out of an app by setting section access, you can open the app without data, and edit the access section in the data load script. This requires that you have access to edit and reload the data load script.

As the same internal logic that is the hallmark of Qlik Sense is also used in the access section, the security fields can be put in different tables. All the fields listed in **LOAD** or **SELECT** statements in the section access must be written in UPPER CASE. Convert any field name containing lower case letters in the database to upper case using the **Upper** function before reading the field by the **LOAD** or **SELECT** statement.

A wildcard, *, is interpreted as all (listed) values of this field, that is. a value listed elsewhere in this table. If used in one of the system fields (USERID, GROUP) in a table loaded in the access section of the script, it is interpreted as all (also not listed) possible values of this field.



When loading data from a QVD file, the use of the upper function will slow down the loading speed.



If you have enabled section access, you cannot use the section access system field names listed here as field names in your data model.

Example:

In this example, only users in the finance group can open the document.

Access to document

Access	Group
USER	Finance

Dynamic data reduction

Qlik Sense supports dynamic data reduction, in which some of the data in an app can be hidden from a user, based on the section access login:

- Fields (columns) can be hidden by using the system field OMIT.
- Records (rows) can be hidden by linking the section access data with the real data: The selection of values to be shown or excluded is controlled by having one or more fields with common names in section access and section application. After user login, Qlik Sense will attempt to match the selections in fields in the section access to any fields in the section application with exactly the same field names (the field names must be written in UPPER CASE). After the selections have been made, Qlik Sense will permanently hide all data excluded by these selections from the user.



All field names used in the transfer described above and all field values in these fields must be upper case, because all field names and field values are, by default, converted to upper case in section access.



If you want to enable reload of the script in a Qlik Management Console task, the INTERNAL\SA_SCHEDULER account user with ADMIN access is required.

Example: Data reduction based on user id

```
section access;
LOAD * inline [
ACCESS, USERID, REDUCTION, OMIT
USER, AD_DOMAIN\ADMIN, *,
USER, AD_DOMAIN\A, 1,
USER, AD_DOMAIN\B, 2, NUM
USER, AD_DOMAIN\C, 3, ALPHA
ADMIN, INTERNAL\SA_SCHEDULER, *,
];
section application;
T1:
LOAD *,
NUM AS REDUCTION;
LOAD
Chr( RecNo()+ord('A')-1) AS ALPHA,
RecNo() AS NUM
AUTOGENERATE 3;
```

The field REDUCTION (upper case) now exists in both section access and section application (all field values are also upper case). The two fields would normally be totally different and separated, but using section access, these fields are linked and the number of records displayed to the user is reduced.

The field OMIT, in section access, defines the fields that should be hidden from the user.

The result will be:

- User ADMIN can see all fields and only those records other users can see in this example when REDUCTION is 1,2, or 3.
- User A can see all fields, but only those records associated to REDUCTION=1.
- User B can see all fields except NUM, and only those records associated to REDUCTION=2.
- User C can see all fields except ALPHA, and only those records associated to REDUCTION=3.

Example: Data reduction based on user groups

```
section access;
LOAD * inline [
ACCESS, USERID, GROUP, REDUCTION, OMIT
USER, *, ADMIN, *,
USER, *, A, 1,
USER, *, B, 2, NUM
USER, *, C, 3, ALPHA
USER, *, GROUP1, 3,
```

```
ADMIN, INTERNAL\SA_SCHEDULER, *, *,
];
section application;
T1:
LOAD *,
NUM AS REDUCTION;
LOAD
Chr( RecNo()+ord('A')-1) AS ALPHA,
RecNo() AS NUM
AUTOGENERATE 3;
```

The result will be:

- Users belonging to the ADMIN group are allowed to see all data and all fields.
- Users belonging to the A group are allowed to see data associated to REDUCTION=1 across all fields.
- Users belonging to the B group are allowed to see data associated to REDUCTION=2, but not in the NUM field
- Users belonging to the C group are allowed to see data associated to REDUCTION=3, but not in the ALPHA field
- Users belonging to the GROUP1 group are allowed to see data associated to REDUCTION=3 across all fields
- The user INTERNAL\SA_SCHEDULER does not belong to any groups but is allowed to see all data in all fields.



*The wildcard, character *, in this row refers only to all values within the section access table. If there are values in the section application that are not available in the REDUCTION field in section access, they will be reduced.*

Inherited access restrictions



This functionality is not available in Kubernetes.

A binary load will cause the access restrictions to be inherited by the new Qlik Sense app.

4.7 Configuring analytic connections in Qlik Sense Desktop

With analytic connections you are able to integrate external analysis with your business discovery. An analytic connection extends the expressions you can use in load scripts and charts by calling an external calculation engine (when you do this, the calculation engine acts as a server-side extension (SSE)). For example, you could create an analytic connection to R, and use statistical expressions when you load the data.

For Qlik Sense Desktop, the configuration must be done in the *Settings.ini* file.

Do the following:

4 Loading data with the data load script

1. Open the file *Settings.ini*.
For Qlik Sense Desktop it is located in `C:/Users/<User ID>/Documents/Qlik/Sense/` or in `C:/Users/AppData/Local/Programs/Qlik/Sense/Engine.`
For Qlik Sense it is found in: `C:/ProgramData/Qlik/Sense/Engine/.`
2. Add the following configuration (note the empty line at the end):
[Settings 7]
SSEPlugin=<PluginConfig>[;<PluginConfig>...]

Where <PluginConfig> is a comma-separated list of configuration elements containing the following:

<EngineName>, <Address>[, <PathToCertificate>, <RequestTimeout>, <ReconnectTimeout>]



After adding new connections or changing existing connections, a restart of Qlik Sense Desktop is required for the changes to take effect.



Note that the server-side extension (SSE) plugin server must be running before you start Qlik Sense, otherwise the connection will not be established.

Qlik open source SSE repositories

The following two Qlik SSE repositories are open source:

- <https://github.com/qlik-oss/server-side-extension>
Contains the SSE protocol, general documentation, and examples written in Python and C++.
- <https://github.com/qlik-oss/sse-r-plugin>
Contains an R-plugin written in C#, only the source code. You must create the plugin before it can be used.

Description of the elements

<EngineName>: Mapping/alias to the plugin that will be used from within the expressions in the app using the plugin functions, for example, *SSEPython* for a Python plugin.

<Address>: colon-separated list with two elements, and

- <Host>: DNS name (or IP-address) of the plugin.
- <Port>: Port on which the plugin listens, typically 50051.

<PathToCertificate>: File system path to folder containing client certificates required for secure communication with the plugin. Optional. If omitted, insecure communication will be invoked. This path just points to the folder where the certificates are located. You have to make sure that they are actually copied to that folder. The names of the three certificate files must be the following: *root_cert.pem*, *sse_client_cert.pem*, *sse_client_key.pem*. Only mutual authentication (server and client authentication) is allowed.

<RequestTimeout>: Integer (seconds). Optional. Default value is 0 (infinite). Timeout for message duration.

4 Loading data with the data load script

<ReconnectTimeout>: Integer (seconds). Optional. Default value is 20 (seconds). Time before the client tries to reconnect to the plugin after the connection to the plugin was lost.

Examples:

- Example where one SSE plugin server is defined: `SSEPlugin=SSEPython,localhost:50051`
- Example where two SSE plugin servers are defined:
`SSEPlugin=SSEPython,localhost:50051;R,localhost:50053`
- Example where one SSE plugin server is defined without certificate path but with timeouts set:
`SSEPlugin=SSEPython,localhost:50051,,0,20`

5 Managing big data with on-demand apps

On-demand apps enable you to load and analyze big data sources in Qlik Sense Enterprise. Trying to analyze an entire big data store at one time is highly inefficient. Nevertheless, to make representative visualizations, all the data must be discoverable. Qlik Sense on-demand apps give users aggregate views of big data stores and allow them to identify and load relevant subsets of the data for detailed analysis.

On-demand apps expand the potential use cases for data discovery, enabling business users to conduct associative analysis on larger data sources. They allow users to first select data they are interested in discovering insights about and then interactively generate an on-demand app with which they can analyze the data with the full Qlik in-memory capabilities.

5.1 On-demand app components


Qlik Sense manages the loading of big data sources with selection apps that provide aggregated views of the big data and also enable a user to zoom in and analyze finer-grained data. Embedded in each selection app are on-demand app navigation links to one or more template apps used as the basis for creating on-demand apps. Properties of the navigation links and template apps enable you to tightly control the shape and volume of data loaded into on-demand apps.

Apps can be generated repeatedly from the template app to track frequently changing data sets. While the data is filtered according to selections made in the selection app, the on-demand app content is dynamically loaded from the underlying data source. The same on-demand app can be generated multiple times to make fresh analyses of the data as they change.



On-demand app generation is controlled by the On-demand app service. The service is disabled by default and must be enabled before selection and template apps can be linked and on-demand apps generated. The On-demand app service is managed in the Qlik Management Console.

Relations between on-demand app components.



5.2 Constructing on-demand apps

Because on-demand selection and template apps require special load scripting, they are usually created by users with experience writing Qlik Sense load scripts. On-demand selection apps, for example, must load data with a modest level of dimension granularity. On-demand template apps contain load scripts with data binding expressions used to formulate the queries made on the data sources.

A selection app can be linked to multiple template apps, and a single template app can be linked to by multiple selection apps. But the template app's data binding expressions must correspond to fields in the selection apps that link to it. For that reason, selection and template apps tend to be created in conjunction with one another and often by the same experienced script writer.



There are sample on-demand selection and template apps included in the Qlik Sense Enterprise installation at ProgramData\Qlik\Examples\OnDemandApp\sample. This functionality is not available in Kubernetes.

Creating navigation links also requires an understanding of the fields in the selection app that have corresponding bindings in the template app. That is because each navigation link requires an expression that computes the total number of detail records. That total represents the aggregate records accessible by way of the selection state in the selection app. To create that expression requires that the user know how to compute the template app's total record count using fields available in the selection app.

Using selection apps to generate on-demand apps does not require a user to understand the load script. Once an on-demand app navigation link has been created, a user can drag that navigation link onto the selection app's **App navigation** bar to create an app navigation point. On-demand apps are then generated from the navigation point.

Navigation points become available for on-demand app generation when the maximum row calculation from the expression in the navigation link comes within the required range. At that point, the user can generate an on-demand app. The user can also make another set of selections and generate additional apps based on those different selections.

Navigation links have a limit on the number of on-demand apps that can be generated from the link. When the maximum number of apps has been generated, the user who is generating apps from the navigation point must delete one of the existing apps before generating a new on-demand app. The maximum number of generated apps applies to the on-demand app navigation link. If one on-demand app navigation point is created from the navigation link, then that navigation point would be able to create up to the maximum number. When multiple navigation points are created from the same navigation link, together those navigation points are limited to the maximum number set for the navigation link.

Navigation links also set a retention time for generated apps. On-demand apps are automatically deleted when their retention period expires.

5.3 Publishing on-demand apps (Windows)

Most users will use on-demand and selection apps after they have been published. When selection apps are published to a stream, users who have proper permission on that stream can use them to make aggregate selections and generate on-demand apps from the navigation points included with the selection apps. As with all published apps, they cannot be changed after they have been published. To add navigation points, for example, the user would have to make a copy of the selection app.

In many cases, users only use generated on-demand apps. Each generated app can be published separately. In fact, the app navigation link can specify that apps generated from it be published to a specific stream automatically. Users would then explore the selected slices of data loaded with those generated on-demand apps on the stream to which the app was published.

5.4 Sharing on-demand apps (Kubernetes)

You can share on-demand selection apps in your personal space. You must first publish the sheets of the selection and template apps, and then share the selection app. The selection app will then be accessible to the users with whom it is shared. Those users can access the navigation points in those apps. For example, they can generate an on-demand app from the **App navigation** bar. You cannot create or share on-demand apps in shared spaces.

5.5 Advantages of on-demand apps

On-demand apps help business users and IT departments derive value from big data environments in numerous ways. On-demand apps:

- Provide users with a "shopping list" experience that enables them to interactively populate their apps with a subset of data such as time period, customer segment, or geography.
- Provide full Qlik Sense functionality on a latent subset that is hosted in memory. In contrast, Direct Discovery, which can also manage large data sources, does not keep all relevant data in memory. With Direct Discovery, measure data resides at the source until execution.

- Enable IT to govern how large an app can be and invoke apps based on data volume or dimensional selections.
- Provide access to non-SQL data sources such as Teradata Aster, MapR, SAP BEx, and the PLACEHOLDER function in SAP HANA.
Performing non-SQL queries is in contrast to Direct Discovery, which can only be used with SQL data sources.
- Allow customizable SQL and load script generation.
- Allow section access in all cases.

5.6 Creating an on-demand selection app

An on-demand selection app provides the means for selecting subsets of large data sets so that the Qlik associative engine can make associations efficiently and effectively. In very large data volume environments, we recommend that you have the selection app load only a modest level of dimension granularity. For example, a selection app whose data is based on sales data aggregated by quarter, region, and product category could use an **SQL SELECT** statement such as the following:

```
SELECT SUM(S.UNIT_COST) AS TOTAL_UNIT_COST, SUM(S.QUANTITY) AS TOTAL_QUANTITY, SUM(S.UNIT_PRICE * S.QUANTITY) AS TOTAL_SALE, SUM( (S.UNIT_PRICE - S.UNIT_COST) * QUANTITY) AS TOTAL_PROFIT, SUM(1) AS TOTAL_LINE_ITEMS, S.REGION, S.YEARQUARTER, S.PRODCAT, FROM SALE_DETAIL S GROUP BY S.REGION, S.YEARQUARTER, S.PRODCAT
```

The on-demand measure expression property is typically based on a computed aggregate result from an SQL GROUP BY query used to load the data. Because the selection app uses a **GROUP BY** query to aggregate the SALE_DETAIL records, an aggregation function--in this case **SUM**--must be used on the measure fields of UNIT_COST, QUANTITY and the computed values for TOTAL_SALE and TOTAL_PROFIT.

The **SUM(1) AS TOTAL_LINE_ITEMS** provides a way to precisely measure the total number of sale line items for every distinct combination of region, quarter, and product category. When creating a link used to produce on-demand apps, a measure expression must be supplied as way to control the number of records loaded into the on-demand apps. In the SALE_DETAIL example, when a user selects multiple product categories, regions, and/or quarters, a sum can be computed for TOTAL_LINE_ITEMS to determine whether or not the selection exceeds the record limit for the on-demand app.



There is a sample on-demand selection app included in the Qlik Sense Enterprise on Windows installation at ProgramData\Qlik\Examples\OnDemandApp\sample. This functionality is not available in Kubernetes.

Record limits are specified when the selection app is linked to a template app to create an app navigation link. Each app navigation link has a record limit. Multiple navigation links can be created from the selection app. Multiple app navigation links are commonly made linking a selection app to different template apps in order to produce multiple views of data.

Individual on-demand app navigation links can be included in a selection app for publication. Once included in the selection app, an app navigation link is used to create one or more app navigation points that make it possible for users of specific sheets to create on-demand apps based on that link's template app.

5.7 Creating an on-demand template app

An on-demand template app is a regular Qlik Sense app with one important difference: its load script contains data binding expressions used to formulate the queries made on the data sources. These data binding expressions are used at on-demand app generation time to read values from the selection state of the selection app and bind them to the template app script so that the generated app is loaded with a user-controlled subset of the data.

The template app typically connects to the same data source as the selection app. The load script of a selection app typically loads aggregated data to reduce data volumes while still offering interactive visualizations of important dimensions and measures. The load script of a template app uses queries that load a controlled subset of more granular data.



An on-demand template app does not load data directly. Attempting to load data from the template app will result in an error. The template app connection must be valid, but to test whether the connection works correctly, you must generate an on-demand app. When an on-demand app is generated, the load script is modified by the On-demand app service to load the selection state of the on-demand selection app. If the on-demand app generates without error, then you know the connection in template app work correctly.



Consider the size of your apps when developing on-demand template apps in Kubernetes. Depending on your deployment, there may be storage limits, or using a large amount of storage may cause your cloud deployment to scale. Contact your system administrator for more information.

Structure of a template app

A template app is linked to a selection app using an on-demand app navigation link. The app navigation link includes properties that control the number of records queried when the on-demand app is loaded with data. The load script of the template app contains data binding expressions that specify which field data from the selection app is used to formulate the queries issued when loading data into the on-demand app.



There is a sample on-demand template app included in the Qlik Sense Enterprise on Windows installation at `ProgramData\Qlik\Examples\OnDemandApp\sample`. This functionality is not available in Kubernetes.

The data binding expressions have the form:

`$(od_FIELDNAME)`

The `od_` prefix is used to bind the selection state of the selection app to the load script of the on-demand app, which is created by copying the template app. The part of the data binding expression that follows the `od_` prefix must be a name that matches a field in the selection app. When the on-demand app is generated, the current

5 Managing big data with on-demand apps

selection state of the selection app is used to obtain the desired values to bind for each field. Each occurrence of a **\$(od_FIELDNAME)** expression in the load script of the newly created on-demand app is replaced with the list of values selected for the corresponding field in the selection state of the selection app.

To be valid SQL syntax, the template app's **SELECT** statement for filtering on multiple values must use an **IN** clause. The recommended practice is to write a subroutine to create the correct **WHERE** clause:

```
SUB ExtendWhere(Name, ValVarName) LET T = Name & '_COLNAME'; LET ColName = $(T); LET Values = $(ValVarName); IF len(Values) > 0 THEN IF len(WHERE_PART) > 0 THEN LET WHERE_PART = '$(WHERE_PART) AND $(ColName) IN ( $(Values) )'; ELSE LET WHERE_PART = ' WHERE $(ColName) IN ( $(Values) )'; ENDIF ENDIF END SUB;
```

Prior to calling the ExtendWhere subroutine, you must set the WHERE_PART variable.

```
SET WHERE_PART = '';
```

If you want the **SQL** statement to have an unchanging **WHERE** clause in addition to the clauses that will be inserted by the selection app, insert a 'WHERE <condition>' in the **SET** statement. Otherwise, use the simple **SET** statement.

The following is an example of how the ExtendedWhere subroutine is used:

```
FOR EACH fldname IN 'ORIGIN', 'DEST', 'YEAR', 'QUARTER', 'ORIGIN_STATE_ABR', 'DEST_STATE_ABR' LET vallist = $(fldname); WHEN (IsNull(vallist)) LET vallist = ''; IF len(vallist) > 0 THEN CALL ExtendWhere('$(fldname)', 'vallist'); ENDIF NEXT fldname
```

The list of field names should match the names of the fields used in each field binding you created.

Once the value list for each field has been built, a **SELECT** statement can be written. For example:

```
SQL SELECT "MKT_ID", "YEAR", "QUARTER", "ORIGIN", "ORIGIN_STATE_ABR", "DEST", "DEST_STATE_ABR", "TICKET_CARRIER", "FARE_CLASS", "PASSENGERS", "DISTANCE", 1 AS "FLIGHT_COUNT" FROM "SAPH7T"."/QT/AIRPORT_FACT" $(WHERE_PART);
```

The $$(WHERE_PART)$ portion of the **SELECT** statement will be expanded to include the **WHERE** clause generated by the execution of the **FOR-NEXT** loop illustrated above. The list of column expressions that follow the **SELECT** keyword should be modified to match your specific database table's columns.



Avoid using the names of fields from the template app model when creating on-demand app binding variables. Variables defined in the script become available in the template app model that is referenced when creating data visualizations. Choosing on-demand app binding variables that do not overlap with fields in the model will prevent unintentional confusion between fields in the template app model and the on-demand app binding variables in the data load script. A good practice is to establish a prefix for on-demand app binding variables. For example, use X_ORIGIN instead of ORIGIN.

Single Sign-On (SSO)

On-demand apps can use single sign-on (SSO) with data sources that support SSO. The engine and the data source must be configured to allow SSO.



This functionality is not available in Kubernetes.

Once the engine and data source have been configured for SSO, the template app must enable SSO by adding the following syntax to the template app script:

```
///!ODAG_SSO
```

The On-Demand App Service parses the script when an on-demand app is generated and each time it is reloaded.

When an on-demand app is loaded with SSO, the identity of the end user is sent to the data source. The end user must have access to the sources used in the template app's data connections. Only data that user has access to in those sources is loaded, even if a larger set of data is selected.



On-demand apps generated from template apps that use single sign-on (SSO) cannot be published.

Reload nodes for template apps

Administrators can control where on-demand apps are reloaded in a multi-node environment by setting load balancing rules on template apps. When a load balancing rule is set, all apps generated from links to the template app will be loaded according to the rule that applies to the template app.



This functionality is not available in Kubernetes.

Binding expressions in on-demand template apps

Data bindings in a template app specify which data from a corresponding selection app is used to formulate the queries issued when loading data into an on-demand app.

The basic form of binding expressions--`$(od_FIELDNAME)`--can be modified to refine selections and to ensure that the template app loads data correctly.



Template apps originally created using the Qlik Sense extension for On-demand App Generation should be changed to use the approach illustrated below for binding a large number of selections from a field.

Binding a large number of selections from a field

The Qlik associative engine limits the number of values that can be concatenated together in a single **LET** statement. To avoid the limitation, an **INLINE LOAD** statement can be used to efficiently construct a comma-separated list of values.

First, you must build a script subroutine to process an **INLINE** table.

```
SUB BuildValueList(VarName, TableName, ColName, QuoteChrNum) IF $(QuoteChrNum) = 0 THEN LET  
LOADEXPR = 'Concat$(ColName),' & chr(39) & ',' & chr(39) & ')' AS CombinedData'; ELSE LET CHREXP = ' chr(' &
```

5 Managing big data with on-demand apps

```
'$(QuoteChrNum)' & ')'; LET LOADEXPR = 'Concat( $(CHREXPR) & $(ColName) & $(CHREXPR)' & ',' & chr(39) & ',' & chr(39) & ') AS CombinedData'; ENDIF _TempTable: LOAD $(LOADEXPR) Resident $(TableName); Let vNoOfRows = NoOfRows(_TempTable); IF $(vNoOfRows) > 0 THEN LET $(VarName) = Peek('CombinedData',0,_TempTable); ENDIF drop table _TempTable; drop table '$(TableName)'; END SUB;
```



Calls to the BuildValueList subroutine must use specific values for the QuoteChrNum parameter. When the field processed by the subroutine is numeric the parameter must be set to 0. For character data, the parameter must be set to 39.

The binding should then be written using an **INLINE** table to create a structure for the field values that will load regardless of the number of values.

```
SET ORIGIN=""; OdagBinding: LOAD * INLINE [ VAL $(odso_ORIGIN){"quote": "", "delimiter": ""} ]; SET ORIGIN_ COLNAME='ORIGIN'; CALL BuildValueList('ORIGIN', 'OdagBinding', 'VAL', 39);
```

The `$(od_ORIGIN) {"quote": "", "delimiter": ""}` expression will be replaced by a list of ORIGIN field values from the selection app, separated by line breaks. If the ORIGIN field contains the three values BOS, JFK, ORD, then the expanded INLINE table looks as follows:

```
SET ORIGIN=""; OdagBinding: LOAD * INLINE [ VAL BOS JFK ORD ]; SET ORIGIN_COLNAME='ORIGIN'; CALL BuildValueList('ORIGIN', 'OdagBinding', 'VAL', 39);
```

The value of the ORIGIN variable following the call to BuildValueList will be:

```
'BOS','JFK','JFK'
```

Binding optional and excluded values

Sometimes it is necessary to create a filter condition in the query for the on-demand template app's script using fields that are not directly selectable in the selection app. For example, sales records might use a region code like REGION_CODE while the selection app uses a more recognizable REGION_NAME field for making region selections. The selection app can have a model that uses a separate table to associate REGION_CODE with REGION_NAME so users can select values from REGION_NAME to control the selection state.

While selecting values of REGION_NAME causes those values to be placed in the selected state, the values of REGION_CODE are only in the optional state, that is, white rather than green. Furthermore, if the design of the selection app's sheets excludes REGION_CODE from its set of filter panes, there is no way to have the bind expression `$(od_REGION_CODE)` in the script of the on-demand app expand to the list of selected regions because the REGION_CODE values will never actually be selected, that is, made green.

To handle this situation, there is additional syntax to more precisely control which selection state values are used in each data binding. The `od_` prefix in the field name portion in every on-demand bind expression can include a combination of letters to denote whether the values to be used in the binding are those taken from the selected state and the optional state. The valid combinations, using the REGION_CODE example, are:

Valid combinations, using the REGION_CODE

Pattern	Expansion
<code>\$(ods_REGION_CODE)</code>	Selected (green) values of REGION_CODE

Pattern	Expansion
\$(odo_REGION_CODE)	Optional (white) values of REGION_CODE
\$(odso_REGION_CODE)	Selected or optional values of REGION_CODE
\$(od_REGION_CODE)	Same as \$(ods_REGION_CODE), only selected (green) values of REGION_CODE

In the case of the on-demand app for sales data example, the following data binding expression ensures that either the selected or optional values of REGION_CODE are included in the REGION_CODE binding:

```
$(odso_REGION_CODE)
```

Binding numeric values

When the data to be bound to the on-demand app consists of numbers instead of strings, it is useful to disable the quote wrapping behavior on the numeric fields. For example, if the sales records include a numeric DAY_OF_WEEK column and you want the user of the selection app to choose arbitrary combinations of DAY_OF_WEEK, you would augment the aggregation query used for loading the selection app to include DAY_OF_WEEK in both the **SELECT** list as well as the **GROUP BY** list. If quotation marks are wrapped around DAY_OF_WEEK values when they are selected, a runtime query error could result if the database does not support automatic type conversion from string to numeric.

To handle this situation, there is a syntax suffix that can be added to the end of the FIELDNAME portion of the bind expression to force the field binding to use the numeric values from the selection app rather than string values. The suffix is `_n` as in the following **WHERE** clause:

```
WHERE DAY_OF_WEEK IN ( $(od_DAY_OF_WEEK_n) );
```

This would expand to:

```
WHERE DAY_OF_WEEK IN ( 1,2 );
```

Requiring a certain number of selections

In some situations, it may be necessary to require that the on-demand app query contain a specific number or range of values for a specific field. For example, if the on-demand app's query contains a **BETWEEN** clause to obtain all sales between a start and end date, the bind expression for the YEARQUARTER field can have a suffix syntax of `[2]` that will require exactly two values be selected for YEARQUARTER, as in:

```
$(od_YEARQUARTER)[2]
```

The on-demand app navigation point on the selection app will remain disabled as long as there are not exactly two values of YEARQUARTER selected. A message will display to indicate that exactly two values of YEARQUARTER must be selected.

Selection quantity constraints create a prerequisite linkage between the selection app and the on-demand app. This is different from bind expressions that do not use quantity constraints. For example, when the template app's script contains a bind expression without a quantity constraint, as in:

```
$(od_MYFIELD)
```

5 Managing big data with on-demand apps

there is no requirement that the selection app contain a field named MYFIELD nor for there to be any selected values of that field if it does exist. If the selection app does not contain a field named MYFIELD or if the user simply neglects to make any selections from it, the on-demand app navigation point can still become enabled when other selections are made to fulfill the record-limit value condition.

If on the other hand, the bind expression is:

```
$(od_MYFIELD)[1+]
```

there are now two requirements placed on the selection app:

- The selection app must contain a field named MYFIELD.
- The user must select at least one value for MYFIELD.

This type of bind expression must be used carefully because it limits which selection apps can be used with the template app. You should not use this quantity constraint on bindings of a template app unless you are certain you want to impose that selection quantity requirement on all selection apps that link to that template app.

To perform the data binding process, the On-demand app service uses a string substitution approach that is insensitive to comments in the script. This means you should not use bind expressions in comments unless you want those comments to contain the list of bound values following app generation.

Other quantity constraints are possible. The following table summarizes the different combinations of selection quantity constraints.

Different combinations of selection quantity constraints

Constraint pattern	Selection requirement
\$(od_YEARQUARTER)[2]	Exactly 2 values of YEARQUARTER must be selected.
\$(od_YEARQUARTER)[2-4]	Between 2 and 4 values of YEARQUARTER must be selected.
\$(od_YEARQUARTER)[2+]	At least 2 values of YEARQUARTER must be selected.
\$(od_YEARQUARTER)[2-]	At most 2 values of YEARQUARTER can be selected.



The check to determine if all the quantity constraints in the template app have been met is performed during the app generation process. If a quantity constraint is violated, the request to generate the app will be rejected and an error message will be displayed.

Changing the value quotation and delimiter characters

When a list of values from a field selected in an On-demand selection app is inserted into the script of a template app, the values are surrounded by single quotation marks and separated by commas. These are the default characters for quotations and delimiters. These values can be changed in syntax appended to binding statement for each field. For example:

```
$(odso_ORIGIN){"quote": "|", "delimiter": ";"}
```

These new values are then used when formulating the list of bind values taken from the selection app. If, for example, the selected values are the first three months of the year, the list would be constructed as:

```
|January| ; |February| ; |March|
```


The default values for the quotation and delimiter characters work for most standard SQL databases. But they might not work for some SQL databases and do not work for many dynamic data sources such as NoSQL and REST. For those sources, you need append this binding expression to change the quotation and delimiter characters.

Bind data containing quotation marks

If the data in any of the fields used for bind variables contain either single or double quotation marks, you need to change the default behavior of the `INLINE` load statement to not interpret quotation marks as value delimiters. Change the delimiter using the format spec to a character that never occurs in the values of the field being bound, as is shown with the `|` character in the following example:

```
OdagBinding:
LOAD * INLINE [
VAL
$(odso_ORIGIN){"quote": "'", "delimiter": "|"}
]
(ansi, txt, delimiter is '|', embedded labels);
```

5.8 Building an on-demand app

An on-demand app loads a subset of the data that is loaded in aggregate form by a selection app. An on-demand app is constructed by linking an on-demand selection app to an on-demand template app. Selection apps and template apps are the fundamental building blocks of on-demand apps.

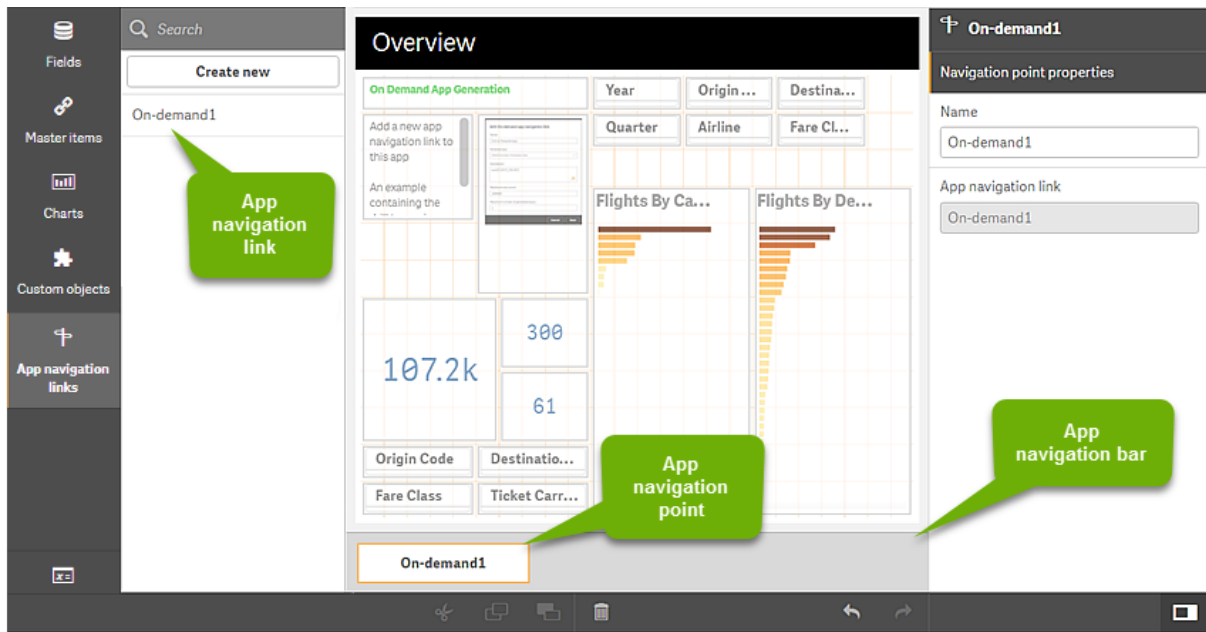
To build an on-demand app, selection and template apps that can be linked together must first be created. To be linked, selection and template apps must have data fields in common that can be bound together.

A selection app can be linked to multiple template apps, and a single template app can be linked to by multiple selection apps. But the template app's data binding expressions must correspond to fields in the selection apps that link to it.

An on-demand app navigation link joins a selection app to a template app. On-demand app navigation links are created in selection apps. Once a navigation link has been defined, it can be added to the selection app's **App navigation** bar as an on-demand app navigation point. Each sheet in an app contains its own **App navigation** bar. Users then generate on-demand apps from the app navigation point.

Example the On-demand app building view.

5 Managing big data with on-demand apps



Multiple on-demand apps, each containing a different combination of selected data, can be generated from the same app navigation point.

Pointers to a single app navigation link can be added to multiple sheets in the same selection app. Also, sheets can have multiple app navigation points, created from multiple app navigation links.

When a selection app is complete with navigation links and navigation points, on-demands can be generated.

Do the following:

1. Open an on-demand selection app and select **Edit**.
2. Select **App navigation links** from the panel on the left side.
3. Click the **Create new** button to open the **Create new On-demand app navigation link** dialog.
4. Name the new on-demand app navigation link.
5. Select an **On-demand template app**.
Not all the apps in the **Template app** drop-down list are valid template apps. You must select an app that has been constructed as a template app and whose data binding expressions correspond to fields in the selection app you are working with. Otherwise, the on-demand apps generated from the app navigation link will produce errors.
6. Write an expression that computes the total number of detail records that are represented by the aggregate records accessible by way of the selection state in the selection app.
The expression usually uses the **SUM** function to obtain a total of the records selected. The result is used to determine when the amount of data to load is within the range specified for generating the on-demand app.
7. Specify the **Maximum row count**.
The **Maximum row count** value sets the upper limit on the number of records, computed by the function in the **Expression** entry, that the on-demand app can load. As long as the number of records as computed by the row estimate expression in the selection app is greater than the **Maximum row count**

value, the on-demand app cannot be generated. The app can only be generated when the number of records computed by the row estimate expression is at or below the upper limit set by the **Maximum row count** value.

To create the expression used for **Maximum row count**, you must know how the total record count is computed from fields available in the selection app.

8. Specify the **Maximum number of generated apps**.

Multiple on-demand apps can be generated from the same on-demand app navigation point on the selection app's **App navigation** bar. The reason for generating multiple apps is that each one can contain a different selection of data. When the maximum number of apps has been generated, the user who is generating apps from the navigation point must delete one of the existing apps before generating a new on-demand app.

The maximum number of generated apps applies to the on-demand app navigation link. If one on-demand app navigation point is created from the navigation link, then that navigation point would be able to create up to the maximum number. But if multiple navigation points are created from the same navigation link, then the total number of on-demand apps generated from those navigation points is limited to the setting for **Maximum number of generated apps**.

9. Enter a numeric value in the **Retention time** field for the length of time apps generated from the navigation link will be retained before they are deleted.

10. In the drop-down menu to the right of the **Retention time** field, select the unit of time for the retention period.

The options for retention time are hours, days, or **Never expires**.

All on-demand apps generated from the navigation link will be retained according to this setting. The age of a generated on-demand app is the difference between the current time and the time of the last data load. This calculation of an on-demand app's age is the same for published and unpublished apps. And if an on-demand app is published manually after it has been generated, the age calculation remains the same: it is based on the last data load of the generated app.



*There is also a retention time setting in the On-Demand App Service that applies to apps generated by anonymous users. That setting does not affect the retention time for users who are logged in with their own identity. For apps generated by anonymous users, the retention time is the shorter of the **Retention time** setting on the navigation link and the On-Demand App Service setting, which is set in the Qlik Management Console. This functionality is not available in Kubernetes.*

11. In the **Default view when opened** drop-down menu, select the sheet to display first when the apps generated from the navigation link are opened.

You can select **App overview** or one of the sheets in the selection app from which the navigation link is created.

12. a. **Windows:** Select a stream from the **Publish to** drop-down menu where apps generated from the navigation link will be published.

You must have permission to publish on the stream you select. If you do not have Publish privileges on the selected stream, attempts to generate on-demand apps from the navigation link will fail.

5 Managing big data with on-demand apps

When selecting a stream to publish generated apps to, you must be sure the intended users of the on-demand app have Read privileges on the stream.

You can also select **Not published (saved to workspace)** to save the generated apps in the users workspace without publishing them.



If anonymous users will be allowed to use a published selection app, the on-demand app navigation links should be configured to publish to a stream that anonymous users can access. If on-demand apps generated from the navigation link are not published automatically, anonymous users will get an error message when they try to generate those apps.

After an app has been generated, it can be published manually.

- b. **Kubernetes:** You can share on-demand selection apps in your personal space. You must first publish the sheets of the selection and template apps, and then share the selection app. The selection app will then be accessible to the users with whom it is shared. Those users can access the navigation points in those apps. For example, they can generate an on-demand app from the **App navigation** bar. You cannot create or share on-demand apps in shared spaces.
13. Click **Create** and the new on-demand app navigation link will appear in the list of **App navigation links**.
14. Drag the app navigation link to the **App navigation** bar on the selection app. Dragging the app navigation link onto the selection app creates an on-demand app navigation point. The properties of the new on-demand app navigation point are displayed in the panel on the right side. You can change the name of the navigation point there if you wish.
15. Click **Done** in the sheet editor. The on-demand selection app is now ready to use or publish. Users of the selection app will be able to generate on-demand apps from the navigation points on the **App navigation** bar in the selection app.

6 Connecting to data sources

Qlik Sense lets you to connect to your data, wherever it is stored, with a wide range of Qlik Connectors and other data connection types. When you create a data connection it is saved to Qlik Sense, so you can quickly select and load data from the data sources that you commonly use. Connect to databases, social media data, local files, remote files, and web files.

6.1 Create a connection

To select data from a data source, you can either create a new data connection or use a saved data connection. You can create data connections and access your saved connections from:

- **Add data** in the data manager.
Add new data to your app quickly and get assistance creating associations.
- **Data connections** in the data load editor.
Select data from a new or existing data connection, or use the script editor to load data from a data connection. You can also edit existing data connections.



You can only see data connections that you own, or have been given access to, for reading or updating. Please contact your Qlik Sense system administrator to acquire access if required.

6.2 Data connection types

Qlik Sense allows you to access your data wherever it resides. The following data connection types are available with Qlik Sense. You can download connectors from the [qlik.com](https://www.qlik.com) download site to use with Qlik Sense.

Many of the connectors that access these data sources are built into Qlik Sense, while others can be added. Each type of data connection has specific settings that you need to configure.

Attached files

Attach data files directly to your app by using drag and drop.

Qlik DataMarket



Qlik Sense Enterprise only.

Select current and historical weather and demographic data, currency exchange rates, as well as business, economic, and societal data.

Database connectors



Qlik Sense Enterprise only.

Connect to an ODBC data source with preconfigured ODBC database connectors.

- Amazon Redshift
- Apache Drill (Beta)
- Apache Hive
- Apache Phoenix (Beta)
- Apache Spark (Beta)
- Azure SQL
- Cloudera Impala
- Google BigQuery
- IBM DB2
- Microsoft SQL Server
- MongoDB (Beta)
- MySQL Enterprise
- Oracle
- PostgreSQL
- Presto
- Sybase ASE
- Teradata

Essbase

Connect to an Essbase dataset.

Local or network files



Qlik Sense Enterprise only.

Connect to a local or network file to select and load data.

ODBC connections through DSN



Qlik Sense Enterprise only.

Connect to a Database Management System (DBMS) with ODBC. Install an ODBC driver for the DBMS in question, and create a data source DSN.

Qlik Associative Big Data Index

The Qlik Associative Big Data Index (QABDI) Selector provides access to a big data source that has been indexed with QABDI.

REST

Connect to a REST data source. The REST connector is not tailored for a specific REST data source and can be used to connect to any data source exposed through the REST API.

Salesforce



Qlik Sense Enterprise only.

Connect to your Salesforce.com account.

SAP



Qlik Sense Enterprise only.

Load data from SAP NetWeaver.

Web files



Qlik Sense Enterprise only

Connect to a web-based data source on a web URL.

Web Storage Provider Connectors



Qlik Sense Enterprise only

Connect to your file-based data from web storage providers.

- Dropbox

Third-party connectors

With third-party connectors, you can connect to data sources that are not directly supported by Qlik Sense. Third-party connectors are developed using the QVX SDK or supplied by third-party developers. In a standard Qlik Sense installation you will not have any third-party connectors available.

6.3 Where is the data connection stored?

Connections are stored in the repository database by the Qlik Sense Repository Service. In a Qlik Sense server deployment, you manage data connections with the Qlik Management Console. The Qlik Management Console allows you to delete data connections, set access rights, and perform other system administration tasks.

In Qlik Sense Desktop, all connections are saved in the app without encryption.



Because Qlik Sense Desktop connections store any details about user name, password, and the file path that you entered when creating the connection, these stored details are available in plain text if you share the app with another user. You need to consider this when you design an app for sharing.

6.4 Limitations

It is not possible to name a data connection 'DM'. This name is reserved by the built-in Qlik DataMarket connector.

6.5 Loading data from files

Qlik Sense can read data from files in a variety of formats:

- Text files, where data in fields is separated by delimiters such as commas, tabs or semicolons (comma-separated variable (CSV) files).
- HTML tables.
- Excel files (except password protected Excel files).
- XML files.
- Qlik native QVD and QVX files.
- Fixed record length files.
- DIF files (Data Interchange Format). DIF files can only be loaded with the data load editor).

You can load files from different data connection types:

- Local and network file folders.
- The **Attached files** folder, which you cannot delete or edit, contains files that are uploaded and attached to the app. (not available in Qlik Sense Desktop).
- Files located on a web URL.

How can I load data from files?

There are several ways to load data from files.

- Adding data with **Add data**, the quickest way to load data from a file. You can load from an existing data connection, or connect to a new data source on the fly.
- Selecting data from a data connection in the data load editor.

Instead of typing the statements manually in the data load editor, you can use the **Select data** dialog to select data to load.

- Loading data from a file by writing script code.
Files are loaded using a **LOAD** statement in the script. **LOAD** statements can include the full set of script expressions.
To read in data from another Qlik Sense app, you can use a **Binary** statement.

Loading files from local and network file folders

You can load files from local and network file folders with a folder connection. You need to make the following settings to the data connection.



This functionality is not available in Kubernetes.

Settings for the data connection

UI item	Description
Path	<p>Path to the folder containing the data files. You can either: Select the folder, type a valid local path or type a UNC path.</p> <p>Example of valid local path: <code>C:\data\MyData\</code></p> <p>Example of a UNC path: <code>\\myserver\filedir\</code></p> <div data-bbox="319 1153 1388 1254"> <i>It is not possible to use a mapped network drive in the path.</i> </div>
Name	Name of the data connection.

Loading files from web resources

You can load files from web resources, such as FTP, HTTP or HTTPS, with a web file data connection. The file can be of any type supported by Qlik Sense. You need to make the following settings to the data connection.



This functionality is not available in Kubernetes.

Settings for a web file data connection

UI item	Description
URL	<p>Full URL to the web file you want to connect to, including the protocol identifier.</p> <p>Example: http://unstats.un.org/unsd/demographic/products/socind/Dec.%202012/1a.xls</p> <p>If you connect to an FTP file you may need to use special characters, for example : or @, in the user name and password part of the URL. In this case you need to replace special characters with a percent character and the ASCII hexadecimal code of the character. For example, you should replace : with '%3a', and @ with '%40'.</p>
Name	Name of the data connection.

The URL set in the web file data connection is static by default, but you can override the URL with the format specification setting **URL is**. This is useful if you need to load data from dynamically created URLs.

Loading data from a dynamically created URL

In this example we want to load forum posts from the first 10 pages of the New to Qlik Sense forum of Qlik Community. The forum page contains 20 posts on each page, and the final parameter of the URL ,start, sets which post to show as the first post of the page. In the example URL here, the page will show posts starting with post number 20, and the following 20 posts.



This functionality is not available in Kubernetes.

<https://community.qlik.com/community/qlik-sense/new-to-qlik-sense/content?filterID=contentstatus%5Bpublished%5D~objecttype~objecttype%5Bthread%5D&itemView=detail&start=20>

With the counter *i* we step through the pages with a step of 20 until 180, which means the **For** loop executes 10 times.

To load the page, we substitute the start page with \$(i) at the end of the URL in the **URL is** setting,.

```
For i = 0 to 180 step 20
```

```
LOAD
```

```
    Title1,
    "Author",
    F6 As Replies,
    Views,
    "Latest activity"
```

```
FROM [lib://x2]
```

```
(URL IS [https://community.qlik.com/community/qlik-sense/new-to-qlik-sense/content?filterID=contentstatus%5Bpublished%5D~objecttype~objecttype%5Bthread%5D&itemView=detail&start=$(i)], html, utf8, embedded labels, table is @1);
```

```
Next i;
```

This will load the 200 most recent posts of the forum in a table, with title, author, number of replies and views, and time of latest activity.

Loading data from Microsoft Excel spreadsheets

Qlik Sense can read data from Microsoft Excel spreadsheets. You can either use Add data in data manager, or select data in the data load editor. In both cases you can select named areas of a sheet, a single sheet, selected sheets, or all sheets from the spreadsheet file. Each sheet is loaded as a separate table, except if they have the same field structure, in which case they are concatenated into one table.



When you load a Microsoft Excel spreadsheet, you are using the spreadsheet as a data source for Qlik Sense apps. That is, Microsoft Excel sheets become tables in Qlik Sense, not sheets in a Qlik Sense app.

You may find it useful to make some changes in Microsoft Excel before you load the spreadsheet.

Selecting data from Microsoft Excel sheets

When you select data from Microsoft Excel sheets, there are some settings to assist you with interpreting the table data correctly:

Settings to assist you with interpreting the table data correctly

UI item	Description
Field names	Set to specify if the table contains Embedded field names or No field names . Typically in an Excel spreadsheet, the first row contains the embedded field names. If you select No field names , fields will be named A,B,C...
Header size	Set to the number of rows to omit as table header, typically rows that contain general information that is not in a columnar format.

Example

My spreadsheet looks like this:

Machine:	AEJ12B	-	-
Date:	2015-10-05 09	-	-
Timestamp	Order	Operator	Yield
2015-10-05 09:22	00122344	A	52
2015-10-05 10:31	00153534	A	67
2015-10-05 13:46	00747899	B	86

In this case you probably want to ignore the two first lines, and load a table with the fields Timestamp, Order, Operator, and Yield. To achieve this, use these settings:

Settings to ignore the two first lines and load the fields

UI item	Description
Header size	2 This means that the first two lines are considered header data and ignored when loading the file. In this case, the two lines starting with Machine: and Date: are ignored, as they are not part of the table data.
Field names	Embedded field names. This means that the first line that is read is used as field names for the respective columns. In this case, the first line that will be read is the third line because the two first lines are header data.

Preparing Microsoft Excel spreadsheets for easier loading with Qlik Sense

If you want to load Microsoft Excel spreadsheets into Qlik Sense, there are many functions you can use to transform and clean your data in the data load script, but it may be more convenient to prepare the source data directly in the Microsoft Excel spreadsheet file. This section provides a few tips to help you prepare your spreadsheet for loading it into Qlik Sense with minimal script coding required.

Use column headings

If you use column headings in Microsoft Excel, they will automatically be used as field names if you select **Embedded field names** when selecting data in Qlik Sense. It is also recommended that you avoid line breaks in the labels, and put the header as the first line of the sheet.

Formatting your data

It is easier to load an Microsoft Excel file into Qlik Sense if the content is arranged as raw data in a table. It is preferable to avoid the following:

- Aggregates, such as sums or counts. Aggregates can be defined and calculated in Qlik Sense.
- Duplicate headers.
- Extra information that is not part of the data, such as comments. The best way is to have a column for comments, that you can easily skip when loading the file in Qlik Sense.
- Cross-table data layout. If, for instance, you have one column per month, you should, instead, have a column called “Month” and write the same data in 12 rows, one row per month. Then you can always view it in cross-table format in Qlik Sense.
- Intermediate headers, for example, a line saying “Department A” followed by the lines pertaining to Department A. Instead, you should create a column called “Department” and fill it with the appropriate department names.
- Merged cells. List the cell value in every cell, instead.
- Blank cells where the value is implied by the previous value above. You need to fill in blanks where there is a repeated value, to make every cell contain a data value.

Use named areas

If you only want to read a part of a sheet, you can select an area of columns and rows and define it as a named area in Microsoft Excel. Qlik Sense can load data from named areas, as well as from sheets.

Typically, you can define the raw data as a named area, and keep all extra commentary and legends outside the named area. This will make it easier to load the data into Qlik Sense.

Remove password protection

Password protected files are not supported by Qlik Sense, so you need to remove password protection from the spreadsheet before loading it into Qlik Sense.

6.6 Loading data from databases

You can load data from commercial database systems into Qlik Sense using the following connectors:

- Connectors specifically developed to load data directly from databases through licensed ODBC drivers, without the need for DSN connections. For more information, see [Qlik Connectors: Database](#).
- Connectors that use the Microsoft ODBC interface or OLE DB. To use Microsoft ODBC, you must install a driver to support your DBMS, and you must configure the database as an ODBC data source in the **ODBC Data Source Administrator** in Windows **Control Panel**.



This functionality is not available in Kubernetes.

Loading data from an ODBC database

There are two ways to load data from a database.

To connect directly to a database through one of the Qlik-licensed ODBC drivers, see the instructions for Database connectors on the Qlik Connectors help site.


For more information, see [Qlik Connectors: Database](#).

Qlik-licensed ODBC drivers support the following databases:

- Amazon Redshift
- Apache Drill (Beta)
- Apache Hive
- Apache Phoenix (Beta)
- Apache Spark (Beta)
- Azure SQL
- Cloudera Impala
- Google BigQuery
- IBM DB2
- Microsoft SQL Server
- MongoDB (Beta)

- MySQL Enterprise
- Oracle
- PostgreSQL
- Presto
- Sybase ASE
- Teradata

To use the Microsoft ODBC interface, do the following:

1. You need to have an ODBC data source for the database you want to access. This is configured in the **ODBC Data Source Administrator** in Windows **Control Panel**. If you do not have one already, you need to add it and configure it, for example pointing to a Microsoft Access database.
2. Open the data load editor.
3. Create an **ODBC** data connection, pointing to the ODBC connection mentioned in step 1.
4. Click  on the data connection to open the data selection dialog.

Now you can select data from the database and insert the script code required to load the data.

ODBC

You can access a DBMS (Database Management System) via ODBC with Qlik Sense:

- You can use the Database connectors in the Qlik ODBC Connector Package that supports the most common ODBC sources. This lets you define the data source in Qlik Sense without the need to use the Microsoft Windows **ODBC Data Source Administrator**. To connect directly to a database through one of the Qlik-licensed ODBC drivers in the ODBC Connector Package, see the instructions for Database connectors on the Qlik Connectors help site.
- You can install an ODBC driver for the DBMS in question, and create a data source DSN. This is described in this section.



This functionality is not available in Kubernetes.



*The **Create new connection (ODBC)** dialog displays the **User DSN** connections that have been configured. When you are using the Qlik Sense Desktop, the list of DSN connections displays the ODBC drivers included in the ODBC Connector Package. They are identified by the "Qlik-" attached to the name (for example, Qlik-db2). These drivers cannot be used to create a new ODBC connection. They are used exclusively by the database connectors in the ODBC Connector Package. The ODBC drivers from the ODBC Connector Package are not displayed when you are using Qlik Sense in a server environment.*

The alternative is to export data from the database into a file that is readable to Qlik Sense.

Normally, some ODBC drivers are installed with Microsoft Windows. Additional drivers can be bought from software retailers, found on the Internet or delivered from the DBMS manufacturer. Some drivers are redistributed freely.

The ODBC interface described here is the interface on the client computer. If the plan is to use ODBC to access a multi-user relational database on a network server, additional DBMS software that allows a client to access the database on the server might be needed. Contact the DBMS supplier for more information on the software needed.

ODBC data connection settings

ODBC data connection settings

UI item	Description
User DSN System DSN	Select which type of DSN to connect to. For User DSN sources you need to specify if a 32-bit driver is used with Use 32-bit connection . System DSN connections can be filtered according to 32-bit or 64-bit .
Single Sign-On	You can enable Single Sign-On (SSO) when connecting to SAP HANA data sources. If this option is not selected, Engine service user credentials are used, unless you specify credentials in Username and Password . If this option is selected, Engine service user or Username / Password credentials are used to do a Windows logon, followed by a subsequent logon to SAML (SAP HANA) using current user credentials.
Username	User name to connect with, if required by the data source. Leave this field empty if you want to use Engine service user credentials, or if the data source does not require credentials.
Password	Password to connect with, if required by the data source. Leave this field empty if you want to use Engine service user credentials, or if the data source does not require credentials.
Name	Name of the data connection.

Adding ODBC drivers

An ODBC driver for your DBMS (DataBase Management System) must be installed for Qlik Sense to be able to access your database. Please refer to the documentation for the DBMS that you are using for further details.

An ODBC driver for your DBMS must be installed for Qlik Sense to be able to access your database. This is external software. Therefore the instructions below may not match the software of all vendors. For details, refer to the documentation for the DBMS you are using.

Do the following:

1. Double-click the **Administrative Tools** icon in the **Control Panel**.
2. Double-click the **Data Sources (ODBC)** icon.
The **ODBC Data Source Administrator** dialog appears.
3. Select the database to use with Qlik Sense.
4. Select the **Drivers** tab in the **Data Sources** dialog.
In the **Drivers** tab you can see a list of installed ODBC drivers. If your DBMS is not listed you must install a driver for it. Run the install program delivered with the ODBC driver, for example the Microsoft ODBC install program.

64-bit and 32-bit versions of ODBC configuration

A 64-bit version of the Microsoft Windows operating system includes the following versions of the Microsoft Open DataBase Connectivity (ODBC) Data Source Administrator tool (*Odbcad32.exe*):

- The 32-bit version of the *Odbcad32.exe* file is located in the `%systemdrive%\Windows\SysWOW64` folder.
- The 64-bit version of the *Odbcad32.exe* file is located in the `%systemdrive%\Windows\System32` folder.

Creating ODBC data sources

An ODBC data source must be created for the database you want to access. This can be done during the ODBC installation or at a later stage.



*Before you start creating data sources, a decision must be made whether the data sources should be **User DSN** or **System DSN** (recommended). You can only reach user data sources with the correct user credentials. On a server installation, typically you need to create system data sources to be able to share the data sources with other users.*

Do the following:

1. Open *Odbcad32.exe*.
2. Go to the tab **System DSN** to create a system data source.
3. Click **Add**.
The **Create New Data Source** dialog appears, showing a list of the ODBC drivers installed.
4. If the correct ODBC driver is listed, select it and click **Finish**.
A dialog specific to the selected database driver appears.
5. Select **Microsoft Access Driver (*.mdb, *.accdb)** and click **Finish**.



If you cannot find this driver in the list you can download it from Microsoft's downloads website and install it.

6. Name the data source and set the necessary parameters.
7. Name the data source *Scripting tutorial ODBC*.
8. Under **Database:**, click **Select....**

9. Under **Directories**, navigate to the location of your *Sales.accdb* file (a tutorial example file).
10. When the file *Sales.accdb* is visible in the text box on the left, click on it to make it the database name.
11. Click **OK** three times to close all the dialogs.
12. Click **OK**.

Best practices when using ODBC data connections

Moving apps with ODBC data connections

If you move an app between Qlik Sense sites/Qlik Sense Desktop installations, data connections are included. If the app contains ODBC data connections, you need to make sure that the related ODBC data sources exist on the new deployment as well. The ODBC data sources need to be named and configured identically, and point to the same databases or files.

Security aspects when connecting to file based ODBC data connections

ODBC data connections using file based drivers will expose the path to the connected data file in the connection string. The path can be exposed when the connection is edited, in the data selection dialog, or in certain SQL queries.

If this is a concern, it is recommended to connect to the data file using a folder data connection if it is possible.

OLE DB

Qlik Sense supports the OLE DB (Object Linking and Embedding, Database) interface for connections to external data sources. A great number of external databases can be accessed via OLE DB.





This functionality is not available in Kubernetes.

OLE DB data connection settings

OLE DB data connection settings

UI item	Description
Provider	Select Provider from the list of available providers. Only available when you create a new connection.

UI item	Description
Data source	<p>Type the name of the Data source to connect to. This can be a server name, or in some cases, the path to a database file. This depends on which OLE DB provider you are using. Only available when you create a new connection.</p> <p>Example:</p> <p>If you selected Microsoft Office 12.0 Access Database Engine OLE DB Provider, enter the file name of the Access database file, including the full file path:</p> <p><i>C:\Users\{user}\Documents\Qlik\Sense\Apps\Tutorial source files\Sales.accdb</i></p> <div style="border: 1px solid gray; padding: 5px; margin-top: 10px;">  <i>If a connection to the data source fails, a warning message is displayed.</i> </div>
Connection string	The connection string to use when connecting to the data source. This string contains references to the Provider and the Data source . Only available when you edit a connection.
Windows integrated security	With this option you use the existing Windows credentials of the user running the Qlik Sense service.
Specific user name and password	With this option you need to enter User name and Password for the data source login credentials.
Username	User name to connect with, if required by the data source. Leave this field empty if you use Windows integrated security or the data source does not require credentials.
Password	Password to connect with, if required by the data source. Leave this field empty if you use Windows integrated security or the data source does not require credentials.
Load Select database...	<p>If you want to test the connection, click Load and then Select database... to use for establishing the data connection.</p> <div style="border: 1px solid gray; padding: 5px; margin-top: 10px;">  <i>You are still able to use all other available databases of the data source when selecting data from the data connection.</i> </div>
Name	Name of the data connection.

Security aspects when connecting to file based OLE DB data connections

OLE DB data connections using file based drivers will expose the path to the connected data file in the connection string. The path can be exposed when the connection is edited, in the data selection dialog, or in certain SQL queries.

If this is a concern, it is recommended that you connect to the data file using a folder data connection if it is possible.

Logic in databases

Several tables from a database application can be included simultaneously in the Qlik Sense logic. When a field exists in more than one table, the tables are logically linked through this key field.

When a value is selected, all values compatible with the selection(s) are displayed as optional. All other values are displayed as excluded.

If values from several fields are selected, a logical AND is assumed.

If several values of the same field are selected, a logical OR is assumed.

In some cases, selections within a field can be set to logical AND.

6.7 Loading data from Qlik DataMarket

You can add data from external sources with Qlik DataMarket. Qlik DataMarket offers an extensive collection of up-to-date and ready-to-use data from external sources accessible directly from within Qlik Sense. Qlik DataMarket provides current and historical weather and demographic data, currency exchange rates, as well as business, economic and societal data.



This functionality is not available in Kubernetes.

Qlik DataMarket also offers data sets from the Eurostat database, including Database by themes, Tables by themes, Tables on EU policy, and Cross cutting topics.

Some Qlik DataMarket data is available for free. Data packages marked **Premium** are available for a subscription fee.

Before you can use Qlik DataMarket data, you must accept the terms and conditions for its use. Also, if you have purchased a license for premium data packages, you must enter your access credentials to use data in those packages. Once access credentials have been applied, the premium data is labeled **Licensed**.

If you accept the terms and conditions but do not enter a license for any of the premium data packages, the premium packages have a **Purchase** button next to them that enables you to buy a license. The **Purchase** button replaces the **Premium** label.

For more information, see [Activating the Qlik DataMarket license](#).



It is not necessary to accept Qlik DataMarket terms and conditions when using Qlik Sense Desktop. Access credentials are also not required because the premium data sets are not available on Qlik Sense Desktop.



The DataMarket user interface can be blocked by browser extensions, such as Privacy Badger, that block ads and enhance privacy. This occurs if the extension mistakes DataMarket's communications for user-tracking by a third party. If you encounter this, you can access DataMarket by excluding your Qlik Sense site from the list of blocked sites in the browser extension that blocks DataMarket.

Qlik DataMarket data can be examined separately or integrated with your own data. Augmenting internal data with Qlik DataMarket can often lead to richer discoveries..

Qlik DataMarket data is current with the source from which it is derived. The frequency with which source data is updated varies. Weather and market data is typically updated at least once a day, while public population statistics are usually updated annually. Most macro-economic indicators, such as unemployment, price indexes and trade, are published monthly. All updates usually become available in Qlik DataMarket within the same day.

Data selections in Qlik Sense are persistent so that the latest available data is loaded from Qlik DataMarket whenever the data model is reloaded.

Most Qlik DataMarket data is both global and country-specific. For example, world population data is available for 200+ countries and territories. In addition, Qlik DataMarket provides various data for states and regions within the United States and European countries.

6.8 Accessing large data sets with Direct Discovery

Direct Discovery enables you to load big data sets from certain SQL sources that have simple star schema structures and combine them with in-memory data. Selections can be made on in-memory and Direct Discovery data to see associations across the data sets with the Qlik Sense association colors--green, white, and gray. Certain visualizations can analyze data from both data sets together, though there are a number of limitations with this approach. It is not designed to be a real-time solution.



This functionality is not available in Kubernetes.



No new development of Direct Discovery will be undertaken to overcome the limitations.

Qlik Sense on-demand apps provide a more flexible approach to loading and analyzing big data sources.

expands the associative capabilities of the Qlik Sense in-memory data model by providing access to additional source data through an aggregated query that seamlessly associates larger data sets with in-memory data. Direct Discovery enhances business users' ability to conduct associative analysis on big data sources without

limitations. Selections can be made on in-memory and Direct Discovery data to see associations across the data sets with the same Qlik Sense association colors - green, white, and gray. Visualizations can analyze data from both data sets together.


Data is selected for Direct Discovery using a special script syntax, **DIRECT QUERY**. Once the Direct Discovery structure is established, Direct Discovery fields can be used along with in-memory data to create Qlik Sense objects. When a Direct Discovery field is used in a Qlik Sense object, an SQL query is run automatically on the external data source.

On-demand apps provide another method for accessing large data sets. In contrast to Direct Discovery, on-demand apps provide full Qlik Sense functionality on a latent subset that is hosted in memory.


Differences between Direct Discovery and in-memory data

In-memory model

In the Qlik Sense in-memory model, all unique values in the fields selected from a table in the load script are loaded into field structures, and the associative data is simultaneously loaded into the table. The field data and the associative data is all held in memory.




A second, related table loaded into memory would share a common field, and that table might add new unique values to the common field, or it might share existing values.




Direct Discovery

When table fields are loaded with a Direct Discovery **LOAD** statement (**Direct Query**), a similar table is created with only the **DIMENSION** fields. As with the In-memory fields, the unique values for the **DIMENSION** fields are loaded into memory. But the associations between the fields are left in the database.



MEASURE field values are also left in the database.



Once the Direct Discovery structure is established, the Direct Discovery fields can be used with certain visualization objects, and they can be used for associations with in-memory fields. When a Direct Discovery field is used, Qlik Sense automatically creates the appropriate SQL query to run on the external data. When selections are made, the associated data values of the Direct Discovery fields are used in the WHERE conditions of the database queries.

With each selection, the visualizations with Direct Discovery fields are recalculated, with the calculations taking place in the source database table by executing the SQL query created by Qlik Sense. The calculation condition feature can be used to specify when visualizations should be recalculated. Until the condition is met, Qlik Sense does not send queries to recalculate the visualizations.

Performance differences between in-memory fields and Direct Discovery fields

In-memory processing is always faster than processing in source databases. Performance of Direct Discovery reflects the performance of the system running the database processing the Direct Discovery queries.

It is possible to use standard database and query tuning best practices for Direct Discovery. All of the performance tuning should be done on the source database. Direct Discovery does not provide support for query performance tuning from the Qlik Sense app. It is possible, however, to make asynchronous, parallel calls to the database by using the connection pooling capability. The load script syntax to set up the pooling capability is:

```
SET DirectConnectionMax=10;
```


Qlik Sense caching also improves the overall user experience. See *Caching and Direct Discovery* (page 202) below.

Performance of Direct Discovery with **DIMENSION** fields can also be improved by detaching some of the fields from associations. This is done with the **DETACH** keyword on **DIRECT QUERY**. While detached fields are not queried for associations, they are still part of the filters, speeding up selection times.

While Qlik Sense in-memory fields and Direct Discovery **DIMENSION** fields both hold all their data in memory, the manner in which they are loaded affects the speed of the loads into memory. Qlik Sense in-memory fields keep only one copy of a field value when there are multiple instances of the same value. However, all field data is loaded, and then the duplicate data is sorted out.

DIMENSION fields also store only one copy of a field value, but the duplicate values are sorted out in the database before they are loaded into memory. When you are dealing with large amounts of data, as you usually are when using Direct Discovery, the data is loaded much faster as a **DIRECT QUERY** load than it would be through the **SQL SELECT** load used for in-memory fields.

Differences between data in-memory and database data

DIRECT QUERY is case-sensitive when making associations with in-memory data. Direct Discovery selects data from source databases according to the case-sensitivity of the database fields queried. If the database fields are not case-sensitive, a Direct Discovery query might return data that an in-memory query would not. For example, if the following data exists in a database that is not case-sensitive, a Direct Discovery query of the value "Red" would return all four rows.

Example table

ColumnA	ColumnB
red	one
Red	two
rED	three
RED	four

An in-memory selection of "Red," on the other hand, would return only:

Red two

Qlik Sense normalizes data to an extent that produces matches on selected data that databases would not match. As a result, an in-memory query may produce more matching values than a Direct Discovery query. For example, in the following table, the values for the number "1" vary by the location of spaces around them:

Table with different values
for the number "1" because
of different location of
spaces around them

ColumnA	ColumnB
' 1'	space_before

ColumnA	ColumnB
'1'	no_space
'1 '	space_after
'2'	two

If you select "1" in a **Filter pane** for ColumnA, where the data is in standard Qlik Sense in-memory, the first three rows are associated:

Associated rows

ColumnA	ColumnB
' 1'	space_before
'1'	no_space
'1 '	space_after

If the **Filter pane** contains Direct Discovery data, the selection of "1" might associate only "no_space". The matches returned for Direct Discovery data depend on the database. Some return only "no_space" and some, like SQL Server, return "no_space" and "space_after".

Caching and Direct Discovery

Qlik Sense caching stores selection states of queries and associated query results in memory. As the same types of selections are made, Qlik Sense leverages the query from the cache instead of querying the source data. When a different selection is made, an SQL query is made on the data source. The cached results are shared across users.

Example:

1. User applies initial selection.
SQL is passed through to the underlying data source.
2. User clears selection and applies same selection as initial selection.
Cache result is returned, SQL is not passed through to the underlying data source.
3. User applies different selection.
SQL is passed through to the underlying data source.

A time limit can be set on caching with the **DirectCacheSeconds** system variable. Once the time limit is reached, Qlik Sense clears the cache for the Direct Discovery query results that were generated for the previous selections. Qlik Sense then queries the source data for the selections and recreates the cache for the designated time limit.

The default cache time for Direct Discovery query results is 30 minutes unless the **DirectCacheSeconds** system variable is used.

Direct Discovery field types

Within Direct Discovery, there are three types of data fields: DIMENSION, MEASURE, and DETAIL. The types are set on data fields when the Direct Discovery selection is made using the **Direct Query** statement in the load script.

All Direct Discovery fields can be used in combination with in-memory fields. Typically, fields with discrete values that will be used as dimensions should be loaded with the DIMENSION keyword, whereas numeric data that will be used in aggregations only should be marked as MEASURE fields.

The following table summarizes the characteristics and usage of the Direct Discovery field types:

Direct Discovery field types

Field Type	In memory?	Forms association?	Used in chart expressions?
DIMENSION	Yes	Yes	Yes
MEASURE	No	No	Yes
DETAIL	No	No	No

DIMENSION fields

DIMENSION fields are loaded in memory and can be used to create associations between in-memory data and the data in Direct Discovery fields. Direct Discovery DIMENSION fields are also used to define dimension values in charts.

MEASURE fields

MEASURE fields, on the other hand, are recognized on a "meta level." MEASURE fields are not loaded in memory (they do not appear in the data model viewer). The purpose is to allow aggregations of the data in MEASURE fields to take place in the database rather than in memory. Nevertheless, MEASURE fields can be used in expressions without altering the expression syntax. As a result, the use of Direct Discovery fields from the database is transparent to the end user.

The following aggregation functions can be used with MEASURE fields:

- **Sum**
- **Avg**
- **Count**
- **Min**
- **Max**

DETAIL fields

DETAIL fields provide information or details that you may want to display but not use in chart expressions. Fields designated as DETAIL commonly contain data that cannot be aggregated in any meaningful way, like comments.

Any field can be designated as a DETAIL field.

Data sources supported in Direct Discovery

Qlik Sense Direct Discovery can be used against the following data sources, both with 32-bit and 64-bit connections:

- ODBC/OLEDB data sources - All ODBC/OLEDB sources are supported, including SQL Server, Teradata and Oracle.
- Connectors that support SQL – SAP SQL Connector, Custom QVX connectors for SQL compliant data stores.

Both 32-bit and 64-bit connections are supported.

SAP

For SAP, Direct Discovery can be used only with the Qlik SAP SQL Connector, and it requires the following parameters in **SET** variables:

```
SET DirectFieldColumnDelimiter=' ';  
SET DirectIdentifierQuoteChar=' ';
```

SAP uses OpenSQL, which delimits columns with a space rather than a comma, so the above set statements cause a substitution to accommodate the difference between ANSI SQL and OpenSQL.

Google Big Query

Direct Discovery can be used with Google Big Query and requires the following parameters in the set variables:

```
SET DirectDistinctSupport=false;  
SET DirectIdentifierQuoteChar='[]';  
SET DirectIdentifierQuoteStyle='big query'
```

Google Big Query does not support **SELECT DISTINCT** or quoted column/table names and has non-ANSI quoting configuration using '[]'.

MySQL and Microsoft Access

Direct discovery can be used in conjunction with MySQL and Microsoft Access but may require the following parameters in the set variables due to the quoting characters used in these sources:

```
SET DirectIdentifierQuoteChar='`';
```

DB2, Oracle and PostgreSQL

Direct discovery can be used in conjunction with DB2, Oracle and PostgreSQL but may require the following parameter in the set variables due to the quoting characters used in these sources:

```
SET DirectIdentifierQuoteChar='\"';
```

Sybase and Microsoft SQL Server

Direct discovery can be used in conjunction with Sybase and Sybase and Microsoft SQL Server but may require the following parameter in the set variables due to the quoting characters used in these sources:

```
SET DirectIdentifierQuoteChar='[]';
```

Apache Hive

Direct discovery can be used in conjunction with Apache Hive but may require the following parameter in the set variables due to the quoting characters used in these sources:

```
SET DirectIdentifierQuoteChar=' ';
```

Cloudera Impala

Direct discovery can be used in conjunction with Cloudera Impala but may require the following parameter in the set variables due to the quoting characters used in these sources:

```
SET DirectIdentifierQuoteChar='[]';
```

This parameter is required when using the Cloudera Impala Connector in the Qlik ODBC Connector Package. It may not be required when using ODBC through DSN.

Limitations when using Direct Discovery

Supported data types

Not all data types are supported in Direct Discovery. There may be cases in which specific source data formats need to be defined in Qlik Sense. You define data formats in the load script by using the "SET Direct...Format" syntax. The following example demonstrates how to define the date format of the source database that is used as the source for Direct Discovery:

Example:

```
SET DirectDateFormat='YYYY-MM-DD';
```

There are also two script variables for controlling how the Direct Discovery formats currency values in the generated SQL statements:

```
SET DirectMoneyFormat (default '#.0000')
SET DirectMoneyDecimalSep (default '.')
```

The syntax for these two variables is the same as for **MoneyFormat** and **MoneyDecimalSep**, but there are two important differences in usage:

- This is not a display format, so it should not include currency symbols or thousands separators.
- The default values are not driven by the locale but are tied to the values. (Locale-specific formats include the currency symbol.)

Direct Discovery can support the selection of extended Unicode data by using the SQL standard format for extended character string literals (N'<extended string>') as required by some databases, such as SQL Server. This syntax can be enabled for Direct Discovery with the script variable **DirectUnicodeStrings**. Setting this variable to "true" enables the use of "N" in front of the string literals.

Security

The following behaviors that could affect security best practice should be taken into consideration when using Direct Discovery:

- All of the users using the same app with the Direct Discovery capability use the same connection. Authentication pass-through and credentials-per-user are not supported.
- Section Access is supported in server mode only.
- Section access is not supported with high-cardinality joins.

- It is possible to execute custom SQL statements in the database with a **NATIVE** keyword expression, so the database connection set up in the load script should use an account that has read-only access to the database.
- Direct Discovery has no logging capability, but it is possible to use the ODBC tracing capability.
- It is possible to flood the database with requests from the client.
- It is possible to get detailed error messages from the server log files.

Qlik Sense functionality that is not supported

Because of the interactive and SQL syntax-specific nature of Direct Discovery, several features are not supported:

- Advanced calculations (set analysis, complex expressions)
- Calculated dimensions
- Comparative analysis (alternate state) of the objects that use Direct Discovery fields
- Direct Discovery **MEASURE** and **DETAIL** fields are not supported in smart search.
- Search of Direct Discovery **DETAIL** fields
- Binary load from an application that is accessing a Direct Discovery table.
- Synthetic keys on the Direct Discovery table
- Table naming in a script does not apply to the Direct Discovery table.
- The wild card character * after a **DIRECT QUERY** keyword in the load script

Example: (DIRECT QUERY *)

- Oracle database tables with LONG data type columns.
- Big integers in scientific notation, outside the range [-9007199254740990, 9007199254740991]. These may cause rounding errors and undefined behavior.
- Snowflake database schemas
- Data preparation in the Data manager
- Qlik Cloud is not supported.
- Qlik Sense Enterprise on Kubernetes
- Export to Microsoft Excel
- Offline mobile ios app is not supported.
- Advanced Analytics Integration is not supported.
- Extensions and Widgets are not supported.
- Qlik GeoAnalytics is not supported.
- Assigning colors to master dimensions and measures is not supported.
- New visualizations included in Qlik Sense June 2017 and later are not supported:
 - Box plot
 - Distribution plot
 - Histogram
- Non-SQL sources and non-SQL statements (for example, the **PLACEHOLDER** function in SAP HANA).

- The following connectors are not supported:
 - Qlik Salesforce Connector
 - Qlik REST Connector
 - Qlik DataMarket Connector
 - Qlik Web connectors
 - Qlik Connector for use with SAP NetWeaver
- Optimizing the SQL generated by the Direct Discovery queries
- High-cardinality joins in combination with in-memory tables can produce large IN clauses that may exceed the SQL buffer limit of the data source.

Multi-table support in Direct Discovery

You can use Direct Discovery to load more than one table or view using ANSI SQL join functionality. In a single chart, all measures must be derived from the same logical table in Qlik Sense, but this can be a combination of several tables from source linked via join statements. However, you can use dimensions sourced from other tables in the same chart.

For example, you can link the tables loaded with Direct Discovery using either a **Where** clause or a **Join** clause.

- Direct Discovery can be deployed in a single fact/multi-dimension in memory scenario with large datasets.
- Direct Discovery can be used with more than one table which match any of the following criteria:
 - The cardinality of the key field in the join is low.
 - The cardinality of the key field in the join is high, **DirectEnableSubquery** is set to true and all tables have been joined with Direct Discovery.
- Direct Discovery is not suitable for deployment in a Third Normal Form scenario with all tables in Direct Discovery form.

Linking Direct Discovery tables with a **Where** clause

In this example script, we load data from the database AW2012. The tables Product and ProductSubcategory are linked with a **Where** clause using the common ProductSubCategoryID field.

```
Product_Join:
DIRECT QUERY
DIMENSION
  [ProductID],
  [AW2012].[Production].[Product].[Name] as [Product Name],
  [AW2012].[Production].[ProductSubcategory].[Name] as [Sub Category Name],
  Color,
  [AW2012].[Production].[Product].ProductSubcategoryID as [SubcategoryID]
MEASURE
  [ListPrice]
FROM [AW2012].[Production].[Product],
  [AW2012].[Production].[ProductSubcategory]
WHERE [AW2012].[Production].[Product].ProductSubcategoryID =
  [AW2012].[Production].[ProductSubcategory].ProductSubcategoryID ;
```

Linking Direct Discovery tables with **Join On** clauses

You can also use **Join On** clauses to link Direct Discovery tables. In this example statement we join the SalesOrderHeader table to the SalesOrderDetail table via the SalesOrderID field, and also join the Customer table to the SalesOrderHeader table via the Customer ID field.


In this example we create measures from the same logical table, which means we can use them in the same chart. For example, you can create a chart with SubTotal and OrderQty as measures.

```
Sales_Order_Header_Join:
DIRECT QUERY
DIMENSION
    AW2012.Sales.Customer.CustomerID as CustomerID,
    AW2012.Sales.SalesOrderHeader.SalesPersonID as SalesPersonID,
    AW2012.Sales.SalesOrderHeader.SalesOrderID as SalesOrderID,
    ProductID,
    AW2012.Sales.Customer.TerritoryID as TerritoryID,
    OrderDate,
    NATIVE('month([OrderDate])') as OrderMonth,
    NATIVE('year([OrderDate])') as OrderYear
MEASURE
    SubTotal,
    TaxAmt,
    TotalDue,
    OrderQty
DETAIL
    DueDate,
    ShipDate,
    CreditCardApprovalCode,
    PersonID,
    StoreID,
    AccountNumber,
    rowguid,
    ModifiedDate
FROM AW2012.Sales.SalesOrderDetail
    JOIN AW2012.Sales.SalesOrderHeader
    ON (AW2012.Sales.SalesOrderDetail.SalesOrderID =
        AW2012.Sales.SalesOrderHeader.SalesOrderID)
    JOIN AW2012.Sales.Customer
    ON(AW2012.Sales.Customer.CustomerID =
        AW2012.Sales.SalesOrderHeader.CustomerID);
```

Using subqueries with Direct Discovery

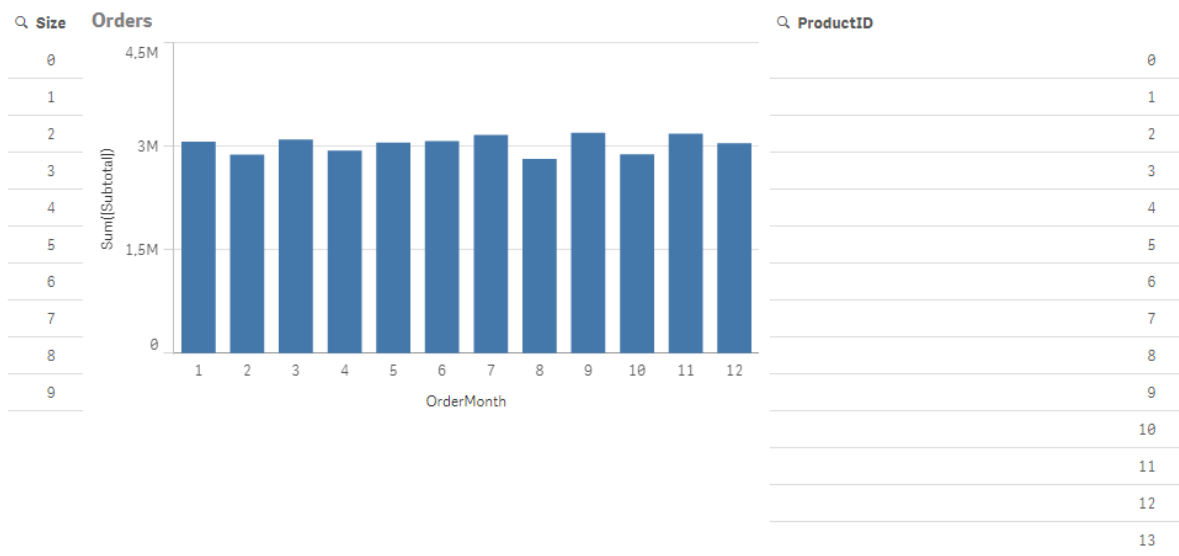
If the cardinality of the key field joining the table is high, that is, it contains a large number of distinct values, a selection in Qlik Sense may generate a very large SQL statement as the **WHERE key_field IN** clause can contain a large number of values. In this case, a possible solution is to let Qlik Sense create subqueries instead.

To illustrate this, we use an example where a products table (ProductTable) is linked to a sales order table (SalesOrderDetail) using a product id (ProductID), with both tables used in Direct Discovery mode.



We create a chart with OrderMonth as dimension, and Sum(Subtotal) as measure, and a filter box for selecting Size.


Orders



Scenario 1: Low cardinality

In this scenario, the product table contains a low number of distinct products, 266. If we make a selection in Size, Direct Discovery generates a SQL statement to return the data, using a **WHERE ProductID IN** clause containing the product IDs matching the selected size, in this case 19 products.

Orders



The SQL that is generated looks like this:

```
SELECT ProductID, month([OrderDate]), SUM(OrderQty), SUM(SubTotal)
FROM SalesTable
WHERE ProductID IN ( 12, 51, 67, 81, 89, 92, 100, 108, 142, 150, 151, 162, 191, 206, 220, 222, 251,
254)
GROUP BY ProductID, month([OrderDate])
```

Scenario 2: Using subqueries

If the same example contains a high number of distinct products, for example 20.000, selecting a dimension filter, Size for example, would generate a SQL statement with a **WHERE ProductID IN** clause containing thousands of product IDs. The resulting statement could be too large to be handled by the data source due to limitations or issues with memory or performance.

The solution is to let Qlik Sense create subqueries instead, by setting the **DirectEnableSubquery** to true. The generated SQL statement could look like this instead:

```
SELECT ProductID, month([OrderDate]), SUM(OrderQty), SUM(SubTotal)
FROM SalesTable
WHERE ProductID IN
( SELECT DISTINCT "AW2012"."dbo"."PRODUCT"."PRODUCTID" WHERE "AW2012"."dbo"."PRODUCT"."SIZE" IN (3) )
GROUP BY ProductID, month([OrderDate])
```

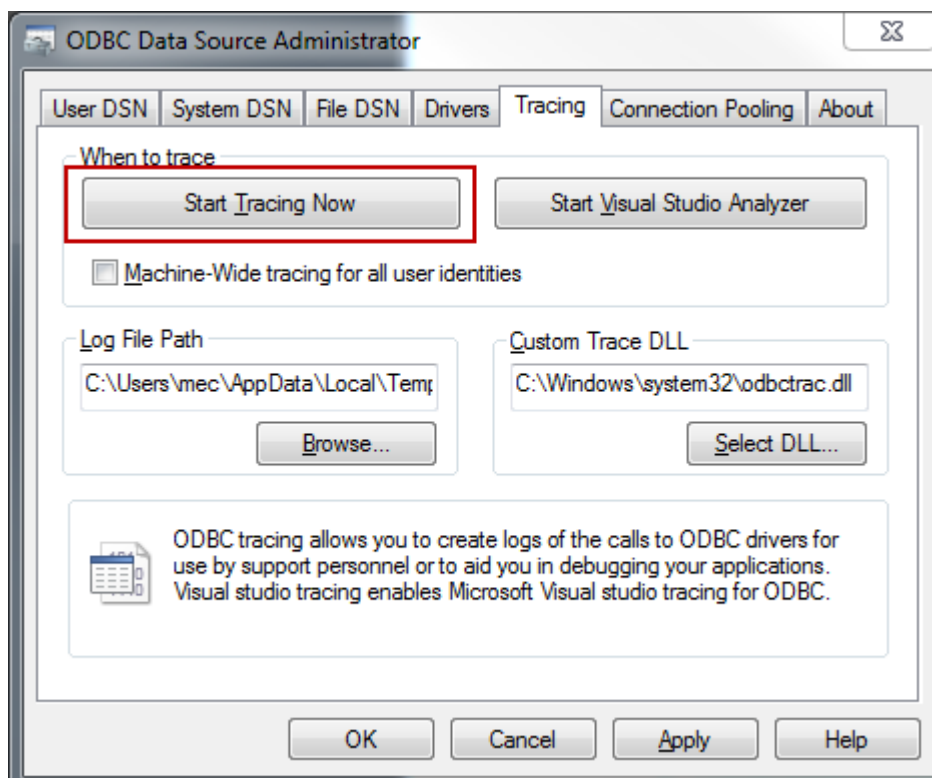
The **WHERE ProductID IN** clause size is not dependent on the number of keys resulting from the selection anymore.

The following limitations apply when using subqueries:

- Subquery syntax is only invoked when you select data which involves filtering a chart using data from another table.
- The amount of data within the keys is the determining factor, not the number of keys.
- Subqueries are only invoked if all tables involved are in Direct Discovery mode. If you filter the chart using data from a table included in memory mode, an **IN** clause will be generated.

Logging Direct Discovery access

Direct Discovery SQL statements passed to the data source can be recorded in the trace files of the database connection. For a standard ODBC connection, tracing is started with the **ODBC Data Source Administrator**:



The resulting trace file details SQL statements generated through the user selections and interactions.

7 Viewing the data model

The data model viewer provides you with an overview of the data structure of the app. You can preview data in the tables and fields in the data model viewer. You can also create dimensions and measures on-the-fly.

The data mode displays the data structure of the app.

The screenshot shows the Qlik Sense data model viewer interface. At the top, there are tabs for 'Data', 'Analysis', and 'Story'. The 'Data' tab is active, displaying a data model with five tables: 'Sales rep', 'Sales', 'Customer', 'Item master', and 'Cities'. Each table is represented by a box containing its name and a list of fields. Lines connect related fields between tables, with a dotted line indicating a circular reference. A green circle 'A' highlights the 'Sales' table, and a green circle 'B' highlights the 'Customer' table. Below the data model, there is a 'Preview' section. On the left, there are buttons for 'Add as dimension' and 'Add as measure'. The 'Preview of data' section shows a table with columns: 'Sales Rep Number', 'Customer Number', 'Item Number', 'IKEY', 'Cost', 'Date', and 'GrossSales'. A green circle 'C' highlights the 'Customer Number' column in the summary table.

Customer Number		Preview of data						
		Sales Rep Number	Customer Number	Item Number	IKEY	Cost	Date	GrossSales
Density	100%	180	10012226	10696	3428	-513.15	1/12/2012	-573.3835
Subset ratio	93.1%	180	10012226	10009	3429	-105.93	1/12/2012	-204.6638
Has duplicates	true	180	10012226	10385	3430	-88.07	1/12/2012	-165.8016
Total distinct values	683	180	10012226	10215	3431	-43.12	1/12/2012	-118.3703
Present distinct values	636	180	10012226	10965	3432	-37.98	1/12/2012	-102.3319
Non-null values	96466	180	10012226	10901	3433	-49.37	1/12/2012	-85.5766
Tags	\$key \$numeric \$integer	180	10012226	10681	3434	-45.81	1/12/2012	-68.4399

In the data model viewer, each data table is represented by a box, with the table name as title and with all fields in the table listed. Table associations are shown with lines, with a dotted line indicating a circular reference. When you select a table or a field, the highlighting of associations instantly gives you a picture of how fields and tables are related.












You can change the zoom level by clicking , or using the slider. Click to restore the zoom level to 1:1.

7.1 Toolbar

In the data model viewer, you find the following tools in the toolbar at the top of the screen:

Toolbar options



UI item	Description
	Global menu with navigation options, and actions that you can perform in your app.

UI item	Description
Data	Click the tab to perform data tasks. For example you can load data in the Data manager or the Data load editor , or view the data model in the Data model viewer . The Data tab is not available in a published app, unless you are the owner of the app. In that case, you can only open the Data model viewer .
Analysis	Click the tab to perform analysis tasks. For example, you can create or interact with tables and charts.
Story	Click the tab to build stories in your app.
Save	Save changes.
	Show or hide app information, where you can choose to edit app information or open app options and style your app.
	Collapse all tables to show the table name only.
	Reduce the size of all tables to show the table name and all fields with associations to other tables.
	Expand all tables to show all fields.
	Internal table view - the Qlik Sense data model including synthetic fields.
	Source table view - the data model of the source data tables.
	Layout menu with the following options:  Grid layout  Auto layout  Restore layout
	Open and close the preview pane.

7.2 Moving and resizing tables in the data model viewer




Moving tables

You can move tables by dragging them on the canvas. The table positions will be saved when the app is saved.

You can lock the table layout (positions and sizes), by clicking  in the right part of the canvas. To unlock the table layout, click .

You can also arrange the layout automatically using the options under  in the toolbar:

Options for moving tables




UI item	Description
 Grid layout	To arrange the tables in a grid.
 Auto layout	To arrange the tables to fit in the window.
 Restore layout	To revert to the layout state present when the data model viewer was last opened.

Resizing tables

You can adjust the display size of a table with the arrow in the bottom right corner of the table. The display size will not be saved when the app is saved.

You can also use the automatic display size options in the toolbar:

Options for resizing tables

UI item	Description
 Collapse all	To minimize all tables to show the table name only
 Show linked fields	To reduce the size of all tables to show the table name and all fields with associations to other tables.
 Expand all	To maximize all tables to show all fields in the table.

7.3 Preview of tables and fields in the data model viewer


In the data model viewer, you can get a preview of any data table in a panel at the bottom of the screen. In the preview, you can quickly inspect the contents of a table or field. You can also add dimensions and measures to the app quickly if you select a field.

Additionally, metadata for the selected table or field are displayed in the preview panel.

You can show and hide the preview panel in two ways:

- Click  in the toolbar.
- Click the **Preview** header.



Direct Discovery data is not displayed in the preview .

Showing a preview of a table

Do the following:

- Click a table header in the data model viewer.

The preview panel is displayed with fields and values of the selected table.

Item master		Preview of data				
Rows	827	Item Number	Product Group	Product Line	Product Sub Group	Product Type
Fields	5	10001	Beverages	Drink	Juice	Pure Juice Beverages
Keys	1	10002	Beverages	Drink	Flavored Drinks	Drinks
Tags	\$key, \$numeric, \$integer\$, \$ascii, \$text	10003	Beverages	Drink	Flavored Drinks	Drinks
		10004	Beverages	Drink	Soda	Carbonated Beverages
		10005	Beverages	Drink	Soda	Carbonated Beverages

Showing a preview of a field

Do the following:

- Click a table field in the data model viewer.

The preview panel is displayed with the selected field and its values, and metadata for the field. You can also add the field as a master dimension or measure.

Product Group		Preview of data			
Density	100%	Item Number	Product Group	Product Line	Product Sub Group
Subset ratio	100%	10001	Beverages	Drink	Juice
Has duplicates	true	10002	Beverages	Drink	Flavored Drinks
Total distinct values	15	10003	Beverages	Drink	Flavored Drinks
Present distinct values	15	10004	Beverages	Drink	Soda
Non-null values	827	10005	Beverages	Drink	Soda
Tags	\$ascii, \$text				

- **Density** is the number of records that have non-NULL values in this field, as compared to the total number of records in the table.
- **Subset ratio** is the number of distinct values of the field found in this table, as compared to the total number of distinct values of this field in other tables in the data model. This is only relevant for key fields.
- If the field is marked with **[Perfect key]**, every row contains a key value that is unique.

7.4 Creating a master dimension from the data model viewer

When you are working with an unpublished app, you can create master dimensions so that they can be reused. Users of a published app will have access to the master dimensions, but will not be able to modify them. The data model viewer is not available in a published app.

Do the following:

1. In the data model viewer, select a field and open the **Preview** panel.
2. Click **Add as dimension**.
The **Create new dimensions** dialog opens, with the selected field. The name of the selected field is also used as the default name of the dimension.
3. Change the name if you want to, and optionally add a description, color, and tags.
4. Click **Add dimension**.
5. Click **Done** to close the dialog.

The dimension is now saved to the master items tab of the assets panel.



*You can quickly add several dimensions as master items by clicking **Add dimension** after adding each dimension. Click **Done** when you have finished.*



Direct Discovery tables are indicated by  in the data model viewer.

7.5 Creating a master measure from the data model viewer

When you are working with an unpublished app, you can create master measures so that they can be reused. Users of a published app will have access to the master measures, but will not be able to modify them.

Do the following:

1. In the data model viewer, select a field and open the **Preview** panel.
2. Click **Add as measure**.
The **Create new measure** dialog opens, with the selected field. The name of the selected field is also used as the default name of the measure.
3. Enter an expression for the measure.
4. Change the name if you want to, and optionally add a description, color, and tags.
5. Click **Create**.

The measure is now saved to the master items tab of the assets panel.

8 Best practices for data modeling

This section describes a number of different ways you can load your data into the Qlik Sense app, depending on how the data is structured and which data model you want to achieve.

8.1 Turning data columns into rows

My data probably looks like this, and I want to have the sales figures in a separate field:

Original data table

Year	Q1	Q2	Q3	Q4
2013	34	54	53	52
2014	47	56	65	67
2015	57	56	63	71

Proposed action

Use the **Crosstable** prefix when you load the table.

The result will look like this:

Table after applying Crosstable prefix

Year	Quarter	Sales
2013	Q1	34
2013	Q2	54
2013	Q3	53
2013	Q4	52
2014	Q1	47
...

8.2 Turning data rows into fields

I have a generic table with three fields similar to this, and I want to have each attribute as a separate table:

Generic table with three fields

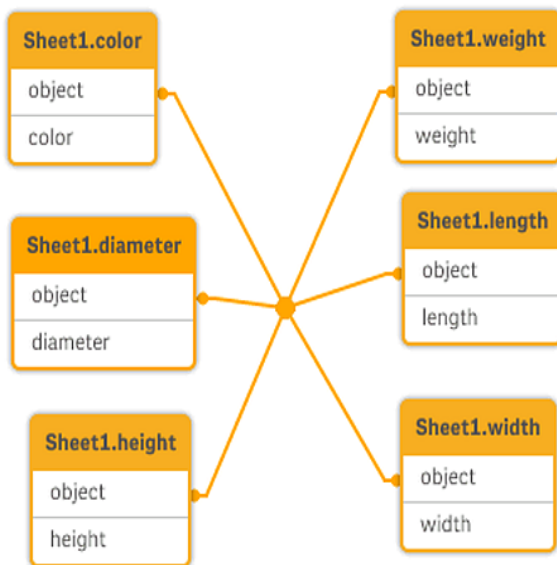
Object	Attribute	Value
ball	color	red
ball	diameter	25
ball	weight	3

Object	Attribute	Value
box	color	56
box	height	30
box	length	20
box	width	25

Proposed action

Create a generic data model using the **Generic** load prefix.

You will get a data model that looks like this:



8.3 Loading data that is organized in hierarchical levels, for example an organization scheme

My data is stored in an adjacent nodes table that looks like this:

Adjacent nodes table

NodeID	ParentNodeID	Title
1	-	General manager
2	1	Country manager
3	2	Region manager

Proposed action

Load the data with the Hierarchy prefix to create an expanded nodes table:

Expanded nodes table

NodeID	ParentNodeID	Title	Level1	Level2	Level3
1	-	General manager	General manager	-	-
2	1	Country manager	General manager	Country manager	-
3	2	Region manager	General manager	Country manager	Region manager

8.4 Loading only new or updated records from a large database

I have a database with a large number of records, and I don't want to reload the entire database to refresh the data in my app. I only want to load new or updated records, and remove records that are deleted from the database.

Proposed action

Implement an incremental load solution using QVD files.

8.5 Combining data from two tables with a common field

Qlik Sense will associate tables with a common field automatically, but I want to control how the tables are combined.

Proposed action : Join / Keep

You can combine two tables into a single internal table with the **Join** or **Keep** prefixes.

Proposed action : Mapping

An alternative to joining tables is to use mapping, which automates lookup of associated values in a mapping table. This can reduce the amount of data to load.

8.6 Matching a discrete value to an interval

I have a table of discrete numeric values (Event), and I want to match it to one or more intervals (Start and End).

Table of discrete numeric values (Event)

Time	Event	Comment
00:00	0	Start of shift 1
01:18	1	Line stop
02:23	2	Line restart 50%
04:15	3	Line speed 100%
08:00	4	Start of shift 2
11:43	5	End of production

8 Best practices for data modeling

Table with intervals (Start and End)

Start	End	Order
01:00	03:35	A
02:30	07:58	B
03:04	10:27	C
07:23	11:43	D

Proposed action

Use the **IntervalMatch** prefix to link the Time field with the interval defined by Start and End.

If the interval is not defined explicitly with start and end, only with a change timestamp like in the table below, you need to create an interval table.

Table with a change timestamp

Currency	Change Data	Rate
EUR	-	8.59
EUR	28/01/2013	8.69
EUR	15/02/2013	8.45
USD	-	6.50
USD	10/01/2013	6.56
USD	03/02/2013	6.30

8.7 Handling inconsistent field values

My data contains field values that are not consistently named in different tables. For example, one table contains the value US in Country while another table contains United States. This situation will prevent associations.

Table 1

Country	Region
US	Maryland
US	Idaho
US	New York
US	California

Table 2

Country	Population
United States	304
Japan	128

Country	Population
Brazil	192
China	1333

Proposed action

Perform data cleansing using a mapping table, that will compare field values and enable correct associations.

8.8 Handling inconsistent field value capitalization

My data contains field values that are not consistently formatted in different tables. For example, one table contains the value single in Type while another table contains Single in the same field. This situation will prevent associations, as the Type field will contain both single and Single values, capitalization matters.

Table 1

Type	Price
single	23
double	39

Table 2

Type	Color
Single	Red
Single	Blue
Double	White
Double	Black

Proposed action

If you loaded the data with **Add data**, you can fix this in the data manager.

Do the following:

1. In the data manager, open Table2 in the table editor.
2. Rename the Type field to Table2.Type.
If you just added the table with **Add data** with data profiling enabled, the field may already be named Table2.Type to prevent automatic association. In this case, this procedure will associate the two tables.
3. Create a calculated field using the expression `Lower(Table2.Type)` and name it Type.
4. Click **Load data**.

Table1 and Table2 should now be associated by the field Type, which only contains values in lowercase, like single and double.

If you want to use different capitalization, you can also achieve this with similar procedures, but remember that the tables will associate using the fields with the same name.

- To get all values capitalized, like Single, create the calculated Type field in Table1 instead, and use the expression `Capitalize(Table1.Type)`.
- To get all values in uppercase, like SINGLE, create the calculated Type field in both tables, and use the expressions `upper(Table1.Type)` and `upper(Table2.Type)` respectively.

8.9 Loading geospatial data to visualize data with a map

I have data that I want to visualize using a map, for example sales data per country, or per store. To use the map visualization I need to load area or point data.

Proposed action

You can load area or point data that match your data value locations from a KML file or an Excel file. Additionally, you need to load the actual map background.

8.10 Loading new and updated records with incremental load

If your app contains a large amount of data from database sources that are continuously updated, reloading the entire data set can be time consuming. In this case you only want to load new or changed records from the database, all other data should already be available in the app. Incremental load, using QVD files, makes it possible to achieve this.

The basic process is described below:

1. Load new or updated data from the database source table.
This is a slow process, but only a limited number of records are loaded.
2. Load data that is already available in the app from the QVD file.
Many records are loaded, but this is a much faster process.
3. Create a new QVD file.
This is the file you will use the next time you do an incremental load.
4. Repeat the procedure for every table loaded.

The following examples show cases where incremental load is used. However, a more complex solution might be necessary, depending on the source database structure and mode of operation.

- Append only (typically used for log files)
- Insert only (no update or delete)
- Insert and update (no delete)
- Insert, update and delete

You can read QVD files in either optimized mode or standard mode. (The method employed is automatically selected by the Qlik Sense engine depending on the complexity of the operation.) Optimized mode is about 10 times faster than standard mode, or about 100 times faster than loading the database in the ordinary fashion.

Append only

The simplest case is the one of log files; files in which records are only appended and never deleted. The following conditions apply:

- The database must be a log file (or some other file in which records are appended and not inserted or deleted) which is contained in a text file (ODBC, OLE DB or other databases are not supported).
- Qlik Sense keeps track of the number of records that have been previously read and loads only records added at the end of the file.

Example:

(Windows)

```
Buffer (Incremental) Load * From LogFile.txt (ansi, txt, delimiter is '\t', embedded labels);
```

Example:

(Kubernetes)

```
Buffer (Incremental) Load * From [lib://MyDataFiles/LogFile.txt] (ansi, txt, delimiter is '\t', embedded labels);
```

Insert only (no update or delete)

If the data resides in a database other than a simple log file, the append approach will not work. However, the problem can still be solved with a minimum amount of extra work. The following conditions apply:

- The data source can be any database.
- Qlik Sense loads records inserted in the database after the last script execution.
- A ModificationTime field (or similar) is required for Qlik Sense to recognize which records are new.

Example:

(Windows)

```
QV_Table:  
SQL SELECT PrimaryKey, X, Y FROM DB_TABLE  
WHERE ModificationTime >= #$(LastExecTime)#  
AND ModificationTime < #$(BeginningThisExecTime)#;
```

```
Concatenate LOAD PrimaryKey, X, Y FROM File.QVD;  
STORE QV_Table INTO File.QVD;
```

The hash signs in the SQL WHERE clause define the beginning and end of a date. Check your database manual for the correct date syntax for your database.

Example:

(Kubernetes)

```
QV_Table:
```

```
SQL SELECT PrimaryKey, X, Y FROM DB_TABLE
WHERE ModificationTime >= #$(LastExecTime)#
AND ModificationTime < #$(BeginningThisExecTime)#;
```

```
Concatenate LOAD PrimaryKey, X, Y FROM [lib://MyDataFiles/File.QVD];
STORE QV_Table INTO [lib://MyDataFiles/File.QVD];
```

The hash signs in the SQL WHERE clause define the beginning and end of a date. Check your database manual for the correct date syntax for your database.

Insert and update (no delete)

The next case is applicable when data in previously loaded records may have changed between script executions. The following conditions apply:

- The data source can be any database.
- Qlik Sense loads records inserted into the database or updated in the database after the last script execution.
- A ModificationTime field (or similar) is required for Qlik Sense to recognize which records are new.
- A primary key field is required for Qlik Sense to sort out updated records from the QVD file.
- This solution will force the reading of the QVD file to standard mode (rather than optimized), which is still considerably faster than loading the entire database.

Example:

(Windows)

```
QV_Table:
SQL SELECT PrimaryKey, X, Y FROM DB_TABLE
WHERE ModificationTime >= #$(LastExecTime)#;

Concatenate LOAD PrimaryKey, X, Y FROM File.QVD
WHERE NOT Exists(PrimaryKey);
WHERE NOT Exists(PrimaryKey);

STORE QV_Table INTO File.QVD;
```

Example:

(Kubernetes)

```
QV_Table:
SQL SELECT PrimaryKey, X, Y FROM DB_TABLE
WHERE ModificationTime >= #$(LastExecTime)#;

WHERE NOT Exists(PrimaryKey);
Concatenate LOAD PrimaryKey, X, Y FROM [lib://MyDataFiles/File.QVD];
WHERE NOT Exists(PrimaryKey);

STORE QV_Table INTO [lib://MyDataFiles/File.QVD];
```


Insert, update and delete

The most difficult case to handle is when records are actually deleted from the source database between script executions. The following conditions apply:

- The data source can be any database.
- Qlik Sense loads records inserted into the database or updated in the database after the last script execution.
- Qlik Sense removes records deleted from the database after the last script execution.
- A field ModificationTime (or similar) is required for Qlik Sense to recognize which records are new.
- A primary key field is required for Qlik Sense to sort out updated records from the QVD file.
- This solution will force the reading of the QVD file to standard mode (rather than optimized), which is still considerably faster than loading the entire database.

Example:

(Windows)

```
Let ThisExecTime = Now( );

QV_Table:
SQL SELECT PrimaryKey, X, Y FROM DB_TABLE
WHERE ModificationTime >= #$(LastExecTime)#
AND ModificationTime < #$(ThisExecTime)#;

Concatenate LOAD PrimaryKey, X, Y FROM File.QVD
WHERE NOT EXISTS(PrimaryKey);

Inner Join SQL SELECT PrimaryKey FROM DB_TABLE;

If ScriptErrorCount = 0 then
STORE QV_Table INTO File.QVD;
Let LastExecTime = ThisExecTime;
End If
```

Example:

(Kubernetes)

```
Let ThisExecTime = Now( );

QV_Table:
SQL SELECT PrimaryKey, X, Y FROM DB_TABLE
WHERE ModificationTime >= #$(LastExecTime)#
AND ModificationTime < #$(ThisExecTime)#;

Concatenate LOAD PrimaryKey, X, Y FROM [lib://MyDataFiles/File.QVD]
WHERE NOT EXISTS(PrimaryKey);

Inner Join SQL SELECT PrimaryKey FROM DB_TABLE;

If ScriptErrorCount = 0 then
STORE QV_Table INTO [lib://MyDataFiles/File.QVD];
```

```
Let LastExecTime = ThisExecTime;  
End If
```

8.11 Combining tables with Join and Keep

A join is an operation that uses two tables and combines them into one. The records of the resulting table are combinations of records in the original tables, usually in such a way that the two records contributing to any given combination in the resulting table have a common value for one or several common fields, a so-called natural join. In Qlik Sense, joins can be made in the script, producing logical tables.

It is possible to join tables already in the script. The Qlik Sense logic will then not see the separate tables, but rather the result of the join, which is a single internal table. In some situations this is needed, but there are disadvantages:

- The loaded tables often become larger, and Qlik Sense works slower.
- Some information may be lost: the frequency (number of records) within the original table may no longer be available.

The **Keep** functionality, which has the effect of reducing one or both of the two tables to the intersection of table data before the tables are stored in Qlik Sense, has been designed to reduce the number of cases where explicit joins need to be used.



In this documentation, the term join is usually used for joins made before the internal tables are created. The association made after the internal tables are created, is however essentially also a join.

Joins within a SQL SELECT statement

With some ODBC drivers it is possible to make a join within the **SELECT** statement. This is almost equivalent to making a join using the **Join** prefix.

However, most ODBC drivers are not able to make a full (bidirectional) outer join. They are only able to make a left or a right outer join. A left (right) outer join only includes combinations where the joining key exists in the left (right) table. A full outer join includes any combination. Qlik Sense automatically makes a full outer join.

Further, making joins in **SELECT** statements is far more complicated than making joins in Qlik Sense.

Example:

```
SELECT DISTINCTROW  
[Order Details].ProductID, [Order Details].  
UnitPrice, Orders.OrderID, Orders.OrderDate, Orders.CustomerID  
FROM Orders  
RIGHT JOIN [Order Details] ON Orders.OrderID = [Order Details].OrderID;
```

This **SELECT** statement joins a table containing orders to a fictive company, with a table containing order details. It is a right outer join, meaning that all the records of *OrderDetails* are included, also the ones with an *OrderID* that does not exist in the table *Orders*. Orders that exist in *Orders* but not in *OrderDetails* are however not included.

Join

The simplest way to make a join is with the **Join** prefix in the script, which joins the internal table with another named table or with the last previously created table. The join will be an outer join, creating all possible combinations of values from the two tables.

Example:

```
LOAD a, b, c from table1.csv;  
join LOAD a, d from table2.csv;
```

The resulting internal table has the fields a, b, c and d. The number of records differs depending on the field values of the two tables.



*The names of the fields to join over must be exactly the same. The number of fields to join over is arbitrary. Usually the tables should have one or a few fields in common. No field in common will render the cartesian product of the tables. All fields in common is also possible, but usually makes no sense. Unless a table name of a previously loaded table is specified in the **Join** statement the **Join** prefix uses the last previously created table. The order of the two statements is thus not arbitrary.*

Keep

The explicit **Join** prefix in the data load script performs a full join of the two tables. The result is one table. In many cases such joins will result in very large tables. One of the main features of Qlik Sense is its ability to make associations between tables instead of joining them, which reduces space in memory, increases speed and gives enormous flexibility. The keep functionality has been designed to reduce the number of cases where explicit joins need to be used.

The **Keep** prefix between two **LOAD** or **SELECT** statements has the effect of reducing one or both of the two tables to the intersection of table data before they are stored in Qlik Sense. The **Keep** prefix must always be preceded by one of the keywords **Inner**, **Left** or **Right**. The selection of records from the tables is made in the same way as in a corresponding join. However, the two tables are not joined and will be stored in Qlik Sense as two separately named tables.

Inner

The **Join** and **Keep** prefixes in the data load script can be preceded by the prefix **Inner**.

If used before **Join**, it specifies that the join between the two tables should be an inner join. The resulting table contains only combinations between the two tables with a full data set from both sides.

If used before **Keep**, it specifies that the two tables should be reduced to their common intersection before being stored in Qlik Sense.

Example:

In these examples we use the source tables Table1 and Table2:

Table 1

A	B
1	aa
2	cc
3	ee

Table2

A	C
1	xx
4	yy

Inner Join

First, we perform an **Inner Join** on the tables, resulting in VTable, containing only one row, the only record existing in both tables, with data combined from both tables.

VTable:

```
SELECT * from Table1;
inner join SELECT * from Table2;
```

VTable

A	B	C
1	aa	xx

Inner Keep

If we perform an **Inner Keep** instead, you will still have two tables. The two tables are of course associated via the common field A.

VTab1:

```
SELECT * from Table1;
```

VTab2:

```
inner keep SELECT * from Table2;
```

VTab1

A	B
1	aa

VTab2

A	C
1	xx

Left

The **Join** and **Keep** prefixes in the data load script can be preceded by the prefix **left**.

8 Best practices for data modeling

If used before **Join**, it specifies that the join between the two tables should be a left join. The resulting table only contains combinations between the two tables with a full data set from the first table.

If used before **Keep**, it specifies that the second table should be reduced to its common intersection with the first table before being stored in Qlik Sense.

Example:

In these examples we use the source tables Table1 and Table2:

Table1

A	B
1	aa
2	cc
3	ee

Table2

A	C
1	xx
4	yy

First, we perform a **Left Join** on the tables, resulting in VTable, containing all rows from Table1, combined with fields from matching rows in Table2.

VTable:

```
SELECT * from Table1;  
left join SELECT * from Table2;
```

VTable

A	B	C
1	aa	xx
2	cc	-
3	ee	-

If we perform an **Left Keep** instead, you will still have two tables. The two tables are of course associated via the common field A.

VTab1:

```
SELECT * from Table1;
```

VTab2:

```
left keep SELECT * from Table2;
```

VTab1

A	B
1	aa
2	cc
3	ee

VTab2

A	C
1	xx

Right

The **Join** and **Keep** prefixes in the data load script can be preceded by the prefix **right**.

If used before **Join**, it specifies that the join between the two tables should be a right join. The resulting table only contains combinations between the two tables with a full data set from the second table.

If used before **Keep**, it specifies that the first table should be reduced to its common intersection with the second table before being stored in Qlik Sense.

Example:

In these examples we use the source tables Table1 and Table2:

Table1

A	B
1	aa
2	cc
3	ee

Table2

A	C
1	xx
4	yy

First, we perform a **Right Join** on the tables, resulting in VTable, containing all rows from Table2, combined with fields from matching rows in Table1.

VTable:

```
SELECT * from Table1;  
right join SELECT * from Table2;
```

VTable

A	B	C
1	aa	xx
4	-	yy

If we perform an **Left Keep** instead, you will still have two tables. The two tables are of course associated via the common field A.

VTab1:

```
SELECT * from Table1;
```

VTab2:

```
right keep SELECT * from Table2;
```

VTab1

A	B
1	aa

VTab2

A	C
1	xx
4	yy

8.12 Using mapping as an alternative to joining

The **Join** prefix in Qlik Sense is a powerful way of combining several data tables in the data model. One disadvantage is that the combined tables can become large and create performance problems. An alternative to **Join** in situations where you need to look up a single value from another table is to use mapping instead. This can save you from loading unnecessary data that slows down calculations and potentially can create calculation errors, as joins can change the number of records in the tables.

A mapping table consists of two columns: a comparison field (input) and a mapping value field (output).

In this example we have a table of orders (Orders), and need to know the countries of the customers, which are stored in the customer table (Customers).

Orders data table

OrderID	OrderDate	ShipperID	Freight	CustomerID
12987	2007-12-01	1	27	3
12988	2007-12-01	1	65	4
12989	2007-12-02	2	32	2
12990	2007-12-03	1	76	3

8 Best practices for data modeling

Customers data table

CustomerID	Name	Country	...
1	DataSales	Spain	...
2	BusinessCorp	Italy	...
3	TechCo	Germany	...
4	Mobecho	France	...

In order to look up the country (Country) of a customer, we need a mapping table that looks like this:

Mapping table

CustomerID	Country
1	Spain
2	Italy
3	Germany
4	France

The mapping table, which we name MapCustomerIDtoCountry, is defined in the script as follows:

```
MapCustomerIDtoCountry:  
Mapping LOAD CustomerID, Country From Customers ;
```

The next step is to apply the mapping, by using the **ApplyMap** function when loading the order table:

Orders:

```
LOAD *,  
    ApplyMap('MapCustomerIDtoCountry', CustomerID, null()) as Country  
From Orders ;
```

The third parameter of the **ApplyMap** function is used to define what to return when a value is not found in the mapping table, in this case **Null()**.

The resulting table will look like this:

Result table

OrderID	OrderDate	ShipperID	Freight	CustomerID	Country
12987	2007-12-01	1	27	3	Germany
12988	2007-12-01	1	65	4	France
12989	2007-12-02	2	32	2	Italy
12990	2007-12-03	1	76	3	Germany

8.13 Working with crosstables in the data load script

A crosstab is a common type of table featuring a matrix of values between two orthogonal lists of header data. It is usually not the optimal data format if you want to associate the data to other data tables.

This topic describes how you can unpivot a crosstab, that is, transpose parts of it into rows, using the **crosstable** prefix to a **LOAD** statement in the data load script.

Unpivoting a crosstab with one qualifying column

A crosstab is often preceded by a number of qualifying columns, which should be read in a straightforward way. In this case there is one qualifying column, Year, and a matrix of sales data per month.

Crosstab with one qualifying column

Year	Jan	Feb	Mar	Apr	May	Jun
2008	45	65	78	12	78	22
2009	11	23	22	22	45	85
2010	65	56	22	79	12	56
2011	45	24	32	78	55	15
2012	45	56	35	78	68	82

If this table is simply loaded into Qlik Sense, the result will be one field for *Year* and one field for each of the months. This is generally not what you would like to have. You would probably prefer to have three fields generated:

- The qualifying column, in this case *Year*, marked with green in the table above.
- The attribute field, in this case represented by the month names Jan - Jun marked with yellow. This field can suitably be named *Month*.
- The data matrix values, marked with blue. In this case they represent sales data, so this can suitably be named *Sales*.

This can be achieved by adding the **crosstable** prefix to the **LOAD** or **SELECT** statement, for example:

```
crosstable (Month, Sales) LOAD * from ex1.xlsx;
```

This creates the following table in Qlik Sense:

Table with crosstable prefix added to the LOAD or SELECT statement

Year	Month	Sales
2008	Jan	45
2008	Feb	65
2008	Mar	78

Year	Month	Sales
2008	Apr	12
2008	May	78
2008	Jun	22
2009	Jan	11
2009	Feb	23
...

Unpivoting a crosstab with two qualifying columns

In this case there are two qualifying columns to the left, followed by the matrix columns.

Crosstab with two qualifying columns

Salesman	Year	Jan	Feb	Mar	Apr	May	Jun
A	2008	45	65	78	12	78	22
A	2009	11	23	22	22	45	85
A	2010	65	56	22	79	12	56
A	2011	45	24	32	78	55	15
A	2012	45	56	35	78	68	82
B	2008	57	77	90	24	90	34
B	2009	23	35	34	34	57	97
B	2010	77	68	34	91	24	68
B	2011	57	36	44	90	67	27
B	2012	57	68	47	90	80	94

The number of qualifying columns can be stated as a third parameter to the **crosstable** prefix as follows:

```
crosstable (Month, Sales, 2) LOAD * from ex2.xlsx;
```

This creates the following result in Qlik Sense:

Table with qualifying columns stated as a third parameter to the crosstable prefix

Salesman	Year	Month	Sales
A	2008	Jan	45
A	2008	Feb	65
A	2008	Mar	78
A	2008	Apr	12

Salesman	Year	Month	Sales
A	2008	May	78
A	2008	Jun	22
A	2009	Jan	11
A	2009	Feb	23
...

8.14 Generic databases

A generic database is a table in which the field names are stored as field values in one column, while the field values are stored in a second. Generic databases are usually used for attributes of different objects.

Look at the example GenericTable below. It is a generic database containing two objects, a ball and a box. Obviously some of the attributes, like color and weight, are common to both the objects, while others, like diameter, height, length and width are not.

GenericTable

object	attribute	value
ball	color	red
ball	diameter	10 cm
ball	weight	100 g
box	color	black
box	height	16 cm
box	length	20 cm
box	weight	500 g
box	width	10 cm

On one hand it would be awkward to store the data in a way giving each attribute a column of its own, since many of the attributes are not relevant for a specific object.


On the other hand, it would look messy displaying it in a way that mixed lengths, colors and weights.

If this database is loaded into Qlik Sense using the standard way and display the data in a table it looks like this:

8 Best practices for data modeling

object ▲	attribute	value
ball	color	red
ball	diameter	10 cm
ball	weight	100 g
box	color	black
box	height	16 cm
box	length	20 cm
box	weight	500 g
box	width	10 cm

However, if the table is loaded as a generic database, column two and three will be split up into different tables, one for each unique value of the second column:



The syntax for doing this is simple:

Example:

```
Generic SELECT* from GenericTable;
```

It does not matter whether a **LOAD** or **SELECT** statement is used to load a generic database.

8.15 Matching intervals to discrete data

The **intervalmatch** prefix to a **LOAD** or **SELECT** statement is used to link discrete numeric values to one or more numeric intervals. This is a very powerful feature which can be used, for example, in production environments as shown in the example below.

Intervalmatch example

Look at the two tables below. The first table shows the start and end of production of different orders. The second table shows some discrete events. How can we associate the discrete events with the orders, so that we know, for example, which orders were affected by the disturbances and which orders were processed by which shifts?

Table OrderLog

Start	End	Order
01:00	03:35	A
02:30	07:58	B
03:04	10:27	C
07:23	11:43	D

Table EventLog

Time	Event	Comment
00:00	0	Start of shift 1
01:18	1	Line stop
02:23	2	Line restart 50%
04:15	3	Line speed 100%
08:00	4	Start of shift 2
11:43	5	End of production

First, load the two tables as usual and then link the field *Time* to the intervals defined by the fields *Start* and *End*:

```
SELECT * from OrderLog;
SELECT * from EventLog;
Intervalmatch (Time) SELECT Start,End from OrderLog;
```

You can now create a table in Qlik Sense as below:

Table with Time field linked to the intervals defined by the Start and End

Time	Event	Comment	Order	Start	End
0:00	0	Start of shift 1	-	-	-

Time	Event	Comment	Order	Start	End
1:18	1	Line stop	A	1:00	3:35
2:23	2	Line restart 50%	A	1:00	3:35
4:15	3	Line speed 100%	B	2:30	7:58
4:15	3	Line speed 100%	C	3:04	10:....
8:00	4	Start of shift 2	C	3:04	10:....
8:00	4	Start of shift 2	D	7:23	11:....
11:43	5	End of production	E	7:23	11:.....

We can now easily see that mainly order *A* was affected by the line stop but that the reduced line speed affected also orders *B* and *C*. Only the orders *C* and *D* were partly handled by *Shift 2*.

Note the following points when using **intervalmatch**:

- Before the **intervalmatch** statement, the field containing the discrete data points (*Time* in the example above) must already have been read into Qlik Sense. The **intervalmatch** statement does not read this field from the database table!
- The table read in the **intervalmatch LOAD** or **SELECT** statement must always contain exactly two fields (*Start* and *End* in the example above). In order to establish a link to other fields you must read the interval fields together with additional fields in a separate **LOAD** or **SELECT** statement (the first **SELECT** statement in the example above).
- The intervals are always closed. That is, the end points are included in the interval. Non-numeric limits render the interval to be disregarded (undefined) while NULL limits extend the interval indefinitely (unlimited).
- The intervals may be overlapping and the discrete values will be linked to all matching intervals.

Using the extended **intervalmatch** syntax to resolve slowly changing dimension problems

The extended **intervalmatch** syntax can be used for handling of the well-known problem of slowly changing dimensions in source data.

Sample script:

```
SET NullInterpret='';
```

```
IntervalTable:
```

```
LOAD Key, ValidFrom, Team
FROM 'lib://dataq/intervalmatch.xlsx' (ooxml, embedded labels, table is IntervalTable);
```

```
Key:
```

```
LOAD
```

```
Key,
```

```
ValidFrom as FirstDate,
```

```
date(if(Key=previous(Key),
```

```
previous(ValidFrom) - 1)) as LastDate,
```

```
Team
```

8 Best practices for data modeling

```
RESIDENT IntervalTable order by Key, ValidFrom desc;
```

```
drop table IntervalTable;
```

```
Transact:
```

```
LOAD Key, Name, Date, Sales
```

```
FROM 'lib://dataqv/intervalmatch.xlsx' (ooxml, embedded labels, table is Transact);
```

```
INNER JOIN intervalmatch (Date,Key) LOAD FirstDate, LastDate, Key RESIDENT Key;
```

The **nullinterpret** statement is only required when reading data from a table file since missing values are defined as empty strings instead of NULL values.

Loading the data from *IntervalTable* would result in the following table:

Table with data loaded from IntervalTable

Key	FirstDate	Team
000110	2011-01-21	Southwest
000110	-	Northwest
000120	-	Northwest
000120	2013-03-05	Southwest
000120	2013-03-05	Northwest
000120	2013-03-05	Southwest

The **nullasvalue** statement allows NULL values to map to the listed fields.

Create *Key*, *FirstDate*, *LastDate*, (attribute fields) by using **previous** and **order by** and thereafter the *IntervalTable* is dropped having been replaced by this key table.

Loading the data from *Transact* would result in the following table:

Table with data loaded from Transact

Key	Name	Date	Sales
000110	Spengler Aaron	2009-08-18	100
000110	Spengler Aaron	2009-12-25	200
000110	Spengler Aaron	2011-02-03	300
000110	Spengler Aaron	2011-05-05	400
000120	Ballard John	2011-06-04	500
000120	Ballard John	2013-01-20	600
000120	Ballard John	2013-03-10	700
000120	Ballard John	2013-03-13	800
000120	Ballard John	2013-09-21	900

The **intervalmatch** statement preceded by the **inner join** replaces the key above with a synthetic key that connects to the *Transact* table resulting in the following table:

Table with the **intervalmatch** statement preceded by the inner join

Key	Team	Name	FirstDate	LastDate	Date	Sales
000110	Northwest	Spengler Aaron	-	2011-01-20	2009-08-18	100
000110	Northwest	Spengler Aaron	-	2011-01-20	2009-12-25	200
000110	Southwest	Spengler Aaron	2011-01-21		2011-02-03	300
000110	Southwest	Spengler Aaron	2011-01-21		2011-05-05	400
000120	Northwest	Ballard John		2013-01-05	2011-06-04	500
000120	Southwest	Ballard John	2013-01-06	2013-03-04	2013-01-20	600
000120	Southwest	Ballard John	2013-03-05		2013-03-10	700
000120	Southwest	Ballard John	2013-03-05		2013-03-13	800
000120	Southwest	Ballard John	2013-03-05		2013-09-21	900

8.16 Creating a date interval from a single date

Sometimes time intervals are not stored explicitly with a beginning and an end. Instead they are implied by only one field – the change timestamp.

It could be as in the table below where you have currency rates for multiple currencies. Each currency rate change is on its own row; each with a new conversion rate. Also, the table contains rows with empty dates corresponding to the initial conversion rate, before the first change was made.

Currency rates

Currency	Change Date	Rate
EUR	-	8.59
EUR	28/01/2013	8.69
EUR	15/02/2013	8.45
USD	-	6.50
USD	10/01/2013	6.56
USD	03/02/2013	6.30

This table defines a set of non-overlapping intervals, where the begin data is called “Change Date” and the end date is defined by the beginning of the following interval. But since the end date isn’t explicitly stored in a column of its own, we need to create such a column, so that the new table will become a list of intervals.

In this script example the table *In_Rates* is created by an inline load. Make sure that the dates in the Change Date column are in the same format as the local date format.


```
In_Rates:
LOAD * Inline [
Currency,Change Date,Rate
EUR,,8.59
EUR,28/01/2013,8.69
EUR,15/02/2013,8.45
USD,,6.50
USD,10/01/2013,6.56
USD,03/02/2013,6.30
];
```

Do the following:

1. Determine which time range you want to work with. The beginning of the range must be before the first date in the data and the end of the range must be after the last.

```
Let vBeginTime = Num('1/1/2013');
Let vEndTime = Num('1/3/2013');
Let vEpsilon = Pow(2,-27);
```

2. Load the source data, but change empty dates to the beginning of the range defined in the previous bullet. The change date should be loaded as “From Date”. Sort the table first according to Currency, then according to the “From Date” descending so that you have the latest dates on top.

```
Tmp_Rates:
LOAD Currency, Rate,
      Date(If(IsNum([Change Date]), [Change Date], $(#vBeginTime))) as FromDate
Resident In_Rates;
```

3. Run a second pass through data where you calculate the “To Date”. If the current record has a different currency from the previous record, then it is the first record of a new currency (but its last interval), so you should use the end of the range defined in step 1. If it is the same Currency, you should take the “From Date” from the previous record, subtract a small amount of time, and use this value as “To Date” in the current record.

```
Rates:
LOAD Currency, Rate, FromDate,
      Date(If( Currency=Peek('Currency'),
              Peek('FromDate') - $(#vEpsilon),
              $(#vEndTime)
            )) as ToDate
Resident Tmp_Rates
Order By Currency, FromDate Desc;
```

4. Drop the input table and the temporary table.
Drop Table Tmp_Rates;

The script listed below will update the source table in the following manner:

Updated source table

Currency	Rate	FromDate	ToDate
EUR	8.45	15/02/2013	vEndTime
EUR	8.69	28/01/2013	14/02/2013 23:59:59
EUR	8.59	vBeginTime	28/01/2013 23:59:99

8 Best practices for data modeling

Currency	Rate	FromDate	ToDate
USD	6.30	03/02/2013	vEndTime
USD	6.56	10/01/2013	2/02/2013 23:59:59
USD	6.50	vBeginTime	9/01/2013 23:59:59

When the script is run, you will have a table listing the intervals correctly. Use the **Preview** section of the data model viewer to view the resulting table.

Preview of data

Currency	Rate	FromDate	ToDate
EUR	8.45	15/02/2013	01/03/2013
EUR	8.69	28/01/2013	14/02/2013
EUR	8.59	01/01/2013	28/01/2013
USD	6.30	03/02/2013	01/03/2013
USD	6.56	10/01/2013	2/02/2013
USD	6.50	01/01/2013	9/01/2013

This table can subsequently be used in a comparison with an existing date using the **Intervalmatch** methods.

Example:

The entire Qlik Sense script looks like this:

```
Let vBeginTime = Num('1/1/2013');
Let vEndTime = Num('1/3/2013');
Let vEpsilon = Pow(2,-27);
```

In_Rates:

```
LOAD * Inline [
Currency,Change Date,Rate
EUR,,8.59
EUR,28/01/2013,8.69
EUR,15/02/2013,8.45
USD,,6.50
USD,10/01/2013,6.56
USD,03/02/2013,6.30
];
```

Tmp_Rates:

```
LOAD Currency, Rate,
      Date(If(ISNum([Change Date]), [Change Date], $(#vBeginTime))) as FromDate
Resident In_Rates;
```

Rates:

```
LOAD Currency, Rate, FromDate,
      Date(If( Currency=Peek('Currency'),
```

```

    Peek('FromDate') - $(#vEpsilon),
    $(#vEndTime)
  )) as ToDate
Resident Tmp_Rates
Order By Currency, FromDate Desc;

```

```
Drop Table Tmp_Rates;
```

8.17 Loading hierarchy data

Unbalanced n -level hierarchies are often used to represent geographical or organizational dimensions in data, among other things. These types of hierarchies are usually stored in an adjacent nodes table, which is in a table where each record corresponds to a node and has a field that contains a reference to the parent node.

Nodes table

NodeID	ParentNodeID	Title
1	-	General manager
2	1	Region manager
3	2	Branch manager
4	3	Department manager

In such a table, the node is stored on one record only but can still have any number of children. The table may of course contain additional fields describing attributes for the nodes.

An adjacent nodes table is optimal for maintenance, but difficult to use in everyday work. Instead, in queries and analysis, other representations are used. The expanded nodes table is one common representation, where each level in the hierarchy is stored in a separate field. The levels in an expanded nodes table can easily be used e.g. in a tree structure. The **hierarchy** keyword can be used in the data load script to transform an adjacent nodes table to an expanded nodes table.

Example:

```

Hierarchy (NodeID, ParentNodeID, Title, 'Manager') LOAD
  NodeID,
  ParentNodeID,
  Title
FROM 'lib://data/hierarchy.txt' (txt, codepage is 1252, embedded labels, delimiter is ',', msq);

```

Expanded nodes table

NodeID	ParentNodeID	Title	Title1	Title2	Title4	Title4
1	-	General manager	General manager	-	-	-
2	1	Region manager	General manager	Region manager	-	-

NodeID	ParentNodeID	Title	Title1	Title2	Title4	Title4
3	2	Branch manager	General manager	Region manager	Branch manager	-
4	3	Department manager	General manager	Region manager	Branch manager	Department manager

A problem with the expanded nodes table is that it is not easy to use the level fields for searches or selections, since prior knowledge about which level to search or select in is needed. An ancestors table is a different representation that solves this problem. This representation is also called a bridge table.

An ancestors table contains one record for every child-ancestor relation found in the data. It contains keys and names for the children as well as for the ancestors. That is, every record describes which node a specific node belongs to. The **hierarchybelongsto** keyword can be used in the data load script to transform an adjacent nodes table to an ancestors table.

8.18 Loading your own map data

To be able to create a map visualization, you need access to geographical data that connects to the data in your app. Qlik Sense can use:

- Name data in fields to place locations in map layers.
- Fields containing geopoints (latitude and longitude) in WGS-84.
- Fields containing geopoints, polygons, or lines from a geographic data source such as a KML file.
- Fields containing geodata in GeoJSON, LineString, or MultiLineString formats.
- Fields containing non-WGS-84 coordinates (when using a custom map as the base map).

When loading map data in **Data manager** with data profiling enabled, the data profiling service will identify country names, city names, and latitude and longitude fields and load the corresponding geometries into new fields. In **Data load editor**, you can optionally combine coordinate fields into a single field for convenience.

Supported name data for fields in a map visualization

The map visualization can use name data in fields to place locations in map layers. The following location types can be used:

- Continent names
- Country names
- ISO alpha 2 country codes
- ISO alpha 3 country codes
- First-level administrative area names. such as a state or province names
- Second-level administrative area names
- Third order administrative area names
- Fourth order administrative area names
- City, village, or other populated place names

- Postal codes or ZIP Codes
- IATA airport codes
- ICAO airport codes



Availability of locations may vary by country. If the named location is not available, use coordinate or area data for the location.

Qlik Sense uses map and location data obtained from recognized field leaders who use accepted methodologies and best practices in marking borders and naming countries within their mappings. Qlik Sense provides flexibility to enable users to integrate their own, separate background maps. If the standard maps do not fit, Qlik Sense offers the option to load customer provided background maps, borders, and areas.

Loading point and area data from a KML file

You can add data from a KML file into your map in **Data manager** and **Data load editor**. By default, all fields are selected in the data selection dialog, even if they do not contain any data. A KML file could contain, for example, area data but no point data. When adding data from a KML file that contains an empty point or area field to Qlik Sense, you can exclude the empty field without running the risk of creating map dimensions without any data.

When adding a field from a KML field to a map layer, if the name field contains meaningful name data, it should be added as the dimension of the layer. The area or point field should then be added as the **Location field**. There will be no difference in how the data is visualized in the layer and the text in the name field will be shown as a tooltip.



If the KML file does not contain point data, line data, or area data, you cannot load data from that file. If the KML file is corrupt, an error message is displayed, and you will not be able to load the data.

Loading map data with data profiling

When you load geographical data using **Add data** in **Data manager** with data profiling enabled, Qlik Sense will try to recognize if your data contains:

- Country and city names from your data
- Geopoint data (latitude, longitude) for a single location, such as a city
- Area data (polygons of geopoints) to represent regions or countries

If successful, a new field containing geographical information is created automatically.



When using **Add data**, data profiling must be enabled. This is the default selection. If you disable data profiling, the geographical data is not detected and the new field containing geographical information is not created.

If cities are recognized during data preparation, the new field contains geoints, and if countries are recognized the new field contains area polygon data. This field is named **<data field>_GeoInfo**. For example, if your data contains a field named **Office** containing city names, a field with geoints named **Office_GeoInfo** is created.



Qlik Sense analyzes a subset of your data to recognize fields containing cities or countries. If the matching is less than 75 percent, a field with geographical information will not be created. If a field is not recognized as geographical data, you can manually change the field type to geographical data.

For more information, see *Changing field types (page 65)*.

Fields with geographical information do not display the geoint or polygon data in the **Associations** preview panel or in the **Tables** view. Instead, the data is indicated generically as *[GEO DATA]*. This improves the speed with which the **Associations** and **Tables** views are displayed. The data is available, however, when you create visualizations in the **Sheet** view.

Loading and formatting point data

You can create a map by using point data (coordinates). Two formats are supported:

- The point data is stored in two fields, one for latitude and one for longitude. You can add the fields to a point layer in the **Latitude** and **Longitude** fields in the point layer. Optionally, you can combine them into a single field. To combine them into a single field:
 - If you used **Add data** with data profiled enabled to load the table, the latitude and longitude fields are recognized, and a geoint field is created automatically.
 - If you loaded the data using the data load script, you can create a single field with point data in *[x, y]* format, using the function `GeomakePoint()`.
For more information, see *Example: Loading point data from separate latitude and longitude columns with the data load script (page 247)*.
- The point data is stored in one field. Each point is specified as an array of *x* and *y* coordinates: *[x, y]*. With geospatial coordinates, this would correspond to *[longitude, latitude]*.
When using this format and loading the data in **Data load editor**, it is recommended that you tag the point data field with `$geoint;`.
For more information: *Example: Loading point data from a single column with the data load script (page 247)*.

In the following examples we assume that the files contain the same data about the location of a company's offices, but in two different formats.

Example: Loading point data from separate latitude and longitude columns with the data load script

The Excel file has the following content for each office:

- Office
- Latitude
- Longitude
- Number of employees

The load script could look as follows:

```
LOAD
    Office,
    Latitude,
    Longitude,
    Employees
FROM 'lib://Maps/Offices.xls'
(biff, embedded labels, table is (Sheet1$));
```

Combine the data in the fields `Latitude` and `Longitude` to define a new field for the points.

Run the script and create a map visualization. Add the point dimension to your map.

You can choose to create the dimension `Location` in the script by adding the following string above the **LOAD** command:

```
LOAD *, GeoMakePoint(Latitude, Longitude) as Location;
```

The function `GeoMakePoint()` joins the longitude and latitude data together.

It is recommended that you tag the field `Office` with `$geoname` so that it is recognized as the name of a geopoint. Add the following lines after the last string in the **LOAD** command:

```
    TAG FIELDS Office WITH $geoname;
```

The complete script then is as follows:

```
LOAD *, GeoMakePoint(Latitude, Longitude) as Location;
LOAD
    Office,
    Latitude,
    Longitude,
    Employees
FROM 'lib://Maps/Offices.xls'
(biff, embedded labels, table is (Sheet1$));

    TAG FIELDS Office WITH $geoname;
```

Run the script and create a map visualization. Add the point dimension to your map.

Example: Loading point data from a single column with the data load script

The Excel file has the following content for each office:

- Office
- Location
- Number of employees

The load script could look as follows:

```
LOAD
  Office,
  Location,
  Employees
FROM 'lib://Maps/Offices.xls'
(biff, embedded labels, table is (Sheet1$));
```

The field `Location` contains the point data and it is recommended to tag the field with `$geopoint` so that it is recognized as a point data field. It is recommended that you tag the field `office` with `$geoname` so that it is recognized as the name of a geopoint. Add the following lines after the last string in the **LOAD** command:

```
    TAG FIELDS Location WITH $geopoint;

    TAG FIELDS Office WITH $geoname;
```

The complete script then looks as follows:

```
LOAD
  Office,
  Location,
  Employees
FROM 'lib://Maps/Offices.xls'
(biff, embedded labels, table is (Sheet1$));

TAG FIELDS Location WITH $geopoint;

TAG FIELDS Office WITH $geoname;
```

Run the script and create a map visualization. Add the point dimension to your map.

8.19 Data cleansing

When loading data from different tables, note that field values denoting the same thing are not always consistently named. Since this lack of consistency is not only annoying, but also hinders associations, the problem needs to be solved. This can be done in an elegant way by creating a mapping table for the comparison of field values.

Mapping tables

Tables loaded via **mapping load** or **mapping select** are treated differently from other tables. They will be stored in a separate area of the memory and used only as mapping tables during script execution. After the script execution they will be automatically dropped.

Rules:

- A mapping table must have two columns, the first one containing the comparison values and the second the desired mapping values.
- The two columns must be named, but the names have no relevance in themselves. The column names have no connection to field names in regular internal tables.

Using a mapping table

When loading several tables listing countries, you may find that one and the same country has several different names. In this example, the U.S.A. are listed as US, U.S., and United States.

To avoid the occurrence of three different records denoting the United States in the concatenated table, create a table similar to that shown and load it as a mapping table.

The entire script should have the following appearance:

```
CountryMap:
Mapping LOAD x,y from MappingTable.txt
(ansi, txt, delimiter is ',', embedded labels);
Map Country using CountryMap;
LOAD Country, City from CountryA.txt
(ansi, txt, delimiter is ',', embedded labels);
LOAD Country, City from CountryB.txt
(ansi, txt, delimiter is ',', embedded labels);
```

The **mapping** statement loads the file *MappingTable.txt* as a mapping table with the label *CountryMap*.

The **map** statement enables mapping of the field *Country* using the previously loaded mapping table *CountryMap*.

The **LOAD** statements load the tables *CountryA* and *CountryB*. These tables, which will be concatenated due to the fact that they have the same set of fields, include the field *Country*, whose field values will be compared with those of the first column of the mapping table. The field values US, U.S., and United States will be found and replaced by the values of the second column of the mapping table, i.e. *USA*.

The automatic mapping is done last in the chain of events that leads up to the field being stored in the Qlik Sense table. For a typical **LOAD** or **SELECT** statement the order of events is roughly as follows:

1. Evaluation of expressions
2. Renaming of fields by as
3. Renaming of fields by alias
4. Qualification of table name, if applicable
5. Mapping of data if field name matches

This means that the mapping is not done every time a field name is encountered as part of an expression but rather when the value is stored under the field name in the Qlik Sense table.

To disable mapping, use the **unmap** statement.

For mapping on expression level, use the **applymap** function.

8 Best practices for data modeling

For mapping on substring level, use the **mapsubstring** function.

9 Troubleshooting - Loading data

This section describes problems that can occur when loading and modeling data in Qlik Sense.

9.1 A data connection stops working after SQL Server is restarted

Possible cause

If you create a data connection to a SQL Server, and then restart the SQL Server, the data connection may stop working, and you are not able to select data. Qlik Sense has lost connection to the SQL Server and was not able to reconnect.

Proposed action

Qlik Sense:

Do the following:

- Close the app, and open it again from the hub.

Qlik Sense Desktop:

Do the following:

1. Close all apps.
2. Restart Qlik Sense Desktop.

9.2 Synthetic keys warning when loading data

If you have loaded several files, you may receive a warning that synthetic keys have been created after loading the data.

Possible cause

If two tables contain more than one common field, Qlik Sense creates a synthetic key to resolve the linking.

Proposed action

In many cases, you do not need to do anything about synthetic keys if the linking is meaningful, but it is a good idea to review the data structure in the data model viewer.

9.3 Circular references warning when loading data

Possible cause

If you have loaded more than two tables, the tables can be associated in such a way that there is more than one path of associations between two fields, causing a loop in the data structure.

Proposed action

9.4 Data selection problems with an OLE DB data source



This functionality is not available in Kubernetes.

Possible cause

If you are not able to select data from an OLE DB data connection, you need to check how the connection is configured.

Proposed action

Do the following:

1. Verify that the connection string is correctly designed.
2. Verify that you are using appropriate credentials to log on.

9.5 Character set problems with non-ANSI encoded data files

You may experience problems with character encoding in non-ANSI encoded data files when using an ODBC data connection.

Possible cause

ODBC data connections do not provide full capabilities for character set encoding.

Proposed action

Do the following:

- If possible, import the data files using a folder data connection, which supports more options for handling character codes. This is probably the best option if you are loading a Microsoft Excel spreadsheet or a text data file.

9.6 Connector is not working

In Qlik Sense Enterprise on Windows, you are trying to create a data connection to a separately installed connector in the data load editor, but the connection fails, or an existing connection is labeled as unknown.

The connector is not properly installed

Possible cause

The connector is not properly installed according to installation instructions. If an app uses a connector on a multi-node site, the connector needs to be installed on all nodes.

Proposed action

Do the following:

- Verify that the connector is installed according to instructions on all nodes of the site.

The connector is not adapted for Qlik Sense

Possible cause

QlikView connectors need to be adapted for Qlik Sense if you want to be able to select data.

Proposed action (if you developed the connector yourself with the QVX SDK)

Do the following:

- You need to adapt the connector for Qlik Sense with an interface to select data.

Proposed action (if the connector was supplied to you)

Do the following:

- Contact the connector supplier to acquire a Qlik Sense adapted connector.

9.7 Data load script is executed without error, but data is not loaded

The script is executed without syntax or load errors, but data is not loaded as expected. A general recommendation is to activate debug to step through the script and examine execution results, but here are some common causes of error.

A statement is not terminated with a semicolon

Possible cause

You have forgotten to terminate a statement with a semicolon.

Proposed action

Do the following:

- Terminate all statements with a semicolon.

Single quote character inside a string

Possible cause

A string contains a single quote character in, for example, a SET variable statement.

Proposed action

Do the following:

- If a string contains a single quote character, it needs to be escaped with an extra single quote.

9.8 Columns are not lining up as expected when selecting data from a fixed record file

Possible cause

The file uses tab characters to pad the columns. Typically, you will see that the field headings do not line up with the expected data if you select **Field breaks** in the select dialog.

In this case, the tab character is usually equivalent to a number of characters.

Proposed action

Do the following:

1. Select **No field names** in **Field names**.
2. Select **Field breaks**.
3. Increase the setting of **Tab size** until you see the columns lining up with the header.
4. Insert field breaks by clicking at the appropriate column positions.
5. Select **Data preview**.
6. Select **Embedded field names** in **Field names**.

The columns are now lined up properly, and each field should have the correct field name.

9.9 Error message "Invalid path" when attaching a file

Possible cause

The file name is too long. Qlik Sense only supports file names up to 171 characters.

Proposed action

Rename the file to a name that contains less than 172 characters.

9.10 Errors when loading an app converted from a QlikView document

In Qlik Sense Enterprise on Windows, you may receive errors when reloading an app that was converted from a QlikView document due to differences between the two products.

Absolute file path references are used in the script

Possible cause

The load script refers to files using absolute paths, which is not supported in Qlik Sense standard mode. Examples of error messages are "Invalid Path" and "LOAD statement only works with lib:// paths in this script mode".

Proposed action

Do the following:

- Replace all file references with **lib://** references to data connections in Qlik Sense.

Unsupported functions or statements are used in the script

Possible cause

If you get a syntax error when running the script in the data load editor, it may be related to using QlikView script statements or functions that are not supported in Qlik Sense.

Proposed action

Do the following:

- Remove the invalid statement or replace it with a valid one.

9.11 Problems connecting to, and loading data from Microsoft Excel files through ODBC

Possible cause

In Qlik Sense Enterprise on Windows, you may encounter problems when setting up an ODBC data connection to a Microsoft Excel file, or loading data from from Microsoft Excel files through an ODBC data connection. This is commonly due to issues with the ODBCDSN configuration in Windows, or problems with the associated ODBC drivers.

Proposed action

Qlik Sense has native support for loading Microsoft Excel files. If possible, replace the ODBC data connection with a folder data connection that connects to the folder containing the Microsoft Excel files.

9.12 Attaching a file by dropping it in **Add data** does not work

You are trying to attach a file by dragging it from **Windows Explorer** and dropping it on the **Add data Attach files** dialog in Qlik Sense, but the file is not uploaded.



This functionality is not available in Kubernetes.

Possible cause

The file is stored in a ZIP archive. It is not possible to attach individual files from a ZIP archive in Qlik Sense, even though the archive appears like a folder in **Windows Explorer**.

Proposed action

Extract the files from the ZIP archive before attaching them.

9.13 Tables with common fields are not automatically associated by field name

You have added two or more tables using **Add data**. The tables have fields with a common field name, but they are not automatically associated.

Possible cause


When you added the tables, you kept the default option to enable data profiling in the **Add data** dialog. This option auto-qualifies all field names that are common between tables. For example, if you add table A and table B with a common field F1 using this option, the field will be named F1 in table A, and B.F1 in table B. This means that the tables are not automatically associated.

Proposed action

Open **Data manager** and select the **Associations** view. Now you can associate the tables based on data profiling recommendations.

When you have associated the tables, click **Load data**.


9.14 Date fields are not recognized as date fields in sheet view

You have fields containing date or timestamp data, but they are not recognized as date fields in sheet view, that is, they are not indicated with  in the assets panel and other field lists.

Data profiling was disabled when the table was added



Possible cause

When you added the tables, you disabled data profiling from **•••** beside the **Add data** button.

With this option, date and timestamp fields that are recognized will function correctly, but they are not indicated with  in the assets panel and other field lists, and expanded property fields are not available.

Proposed action

Open **Data manager** and click **Load data**.

Now, all date and timestamp fields should be indicated with  in the assets panel of sheet view. If they are still not indicated with , the field data is probably using a format that is not recognized as a date.

Date format was not recognized

Possible cause

The input format of the date field was not recognized when the table was loaded. Usually, Qlik Sense recognizes date fields automatically, based on locale settings and common date formats, but in some cases you may need to specify the input format.

Proposed action

Open **Data manager** and edit the table containing the field that was not recognized as a date. The field is most likely indicated with `[]` as a general field. Change the field type to **Date** or **Timestamp**, with an input format that matches the field data.

9.15 Data manager does not show tables in app that contains data

When opening an app created in a Qlik Sense version earlier than 3.0, Data manager shows no tables and a message is displayed that the app contains no data.

Possible cause

The improved data model in Qlik Sense 3.0 and later requires a data reload to complete data profiling and preparation.

Proposed action

Click **Load data** in Data manager. This requires that the app can access the data sources that are used in the app.

9.16 Data manager workflows are broken for all users creating apps on a server

In Qlik Sense Enterprise on Windows, users get errors when trying to use **Add data** or **Load data** in **Data manager**, or when refreshing the app in the browser.

Possible cause

The **Data manager** uses QVD files to cache loaded data. These files are deleted automatically when they are no longer used, but if a large number accumulate, or they become corrupted, they can cause errors.

Proposed action

Delete the folder containing the QVD files. On a Qlik Sense server, the cache is located at:

```
<Qlik Sense shared folder>\Apps\DataPrepAppCache
```

On a Qlik Sense Desktop, the cache is located at:

```
C:\Users\<username>\Documents\Qlik\Sense\Apps\DataPrepAppCache
```

9.17 Running out of disk space

In Qlik Sense Enterprise on Windows, there are several reasons why a system may run low on disk space, and the data manager's method of caching loaded data in QVD files is one possible cause.

Proposed action

Delete the folder containing the QVD files. On a Qlik Sense server, the cache is located at:

```
<Qlik Sense shared folder>\Apps\DataPrepAppCache
```

On a Qlik Sense Desktop, the cache is located at:

```
C:\Users\<username>\Documents\Qlik\Sense\Apps\DataPrepAppCache
```

9.18 Data load editor does not display the script

When the data load editor is opened, the content of the editor is blank, and the script cannot be edited.

Possible cause

The script contains very complex constructions, for example, a large number of nested if statements.

Proposed action

Open the data load editor in safe mode by adding `/debug/dle_safe_mode` to the URL. This will disable syntax highlighting and auto-complete functions, but you should be able to edit and save the script.



*Consider to move the complex parts of the script to a separate text file, and use the **include** variable to inject it into the script at runtime.*